FORSEC

Security in Highly Connected IT Systems

Results of the Bavarian Research Alliance FORSEC

Editors:
Günther Pernul
Guido Schryen
Rolf Schillinger

FAU FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG

Technische Universität München

UNIVERSITÄT PASSAU

UR Universität Regensburg

Günther Pernul, Guido Schryen, Rolf Schillinger (editors):

Security in Highly Connected IT Systems

Results of the Bavarian Research Alliance FORSEC

(01.09.2013 - 31.12.2017)

Security in Highly Connected IT Systems

Results of the Bavarian Research Alliance FORSEC

Editors:

Günther Pernul, Guido Schryen, Rolf Schillinger

Funded by

Bayerisches Staatsministerium für
Bildung und Kultus, Wissenschaft und Kunst

Günther Pernul, Guido Schryen, Rolf Schillinger
Security in Highly Connected IT Systems: Results of the Bavarian Research Alliance FORSEC
ISBN: 978-3-88246-374-3
DOI: 10.5283/epub.36264

Verlag: Universitätsbibliothek Regensburg, Regensburg 2017

Content

List of Figures

List of Tables

Preface

Information and communication technologies (ICT) continue pervading our society and economy. Fostering the exploitation of emerging ICT capabilities is key to achieving a modern society, in which ICT supports and enables progress in substantially important areas and critical infrastructures, including health, mobility, education, production, finance, and public administration. At the same time, this progress is threatened by a large diversity of security and privacy breaches, which represent a severe abuse of ICT enhancements.

Accounting for the urgent need to protect ICT-enabled societal and economic progress, the Bavarian State Ministry of Education, Science and the Arts funded the Bavarian Research Alliance FORSEC – Security in Highly Connected IT Systems with 3.4 million Euros over a period of more than four years (September 2013 until December 2017). The primary goal of the research alliance, which was constituted by four Bavarian universities, was to conduct fundamental research on the protection of modern highly-connected IT systems from a multi-disciplinary perspective, including technological, economic and behavioral aspects.

This book presents a comprehensive overview of the research results and publications that emerged from the multi-disciplinary and multi-organizational research activities of FORSEC. May the achieved research results provide a fertile basis not only for further academic research but also for the development of ICT-oriented applications, products, tools and policies which protect citizens, organizations, companies, and public authorities from emergent security and privacy breaches.

Munich, December 2017

Dr. Ludwig Spaenle

Bavarian State Minister
of Education, Science and the Arts

Foreword

This book reports the final results of the Bavarian Research Alliance *FORSEC – Security in Highly Connected IT Systems*. FORSEC is a joint research alliance of four Bavarian universities (University of Regensburg, University of Passau, Technical University of Munich and Friedrich-Alexander-University of Erlangen-Nuremberg) and has been generously funded by the Bavarian State Ministry of Education, Science and Arts.

The research alliance FORSEC would not have been possible without the work of our participating colleagues, including the Principal Investigators, doctoral students, and student workers, all of whom spent much time in doing collaborative research, writing publications, organizing and attending workshops and conferences over a period of more than four years. We would like to thank them all for making FORSEC a successful research alliance.

Being a research alliance of four universities, ten research groups, and eleven research projects, FORSEC has gone beyond what can be achieved by a set of individual research projects that are unconnected to each other. The nature of a collaborative research endeavor has been implemented by the provision of overall guiding research questions, the organizational union of the research projects to overall three research clusters, the conducting of several workshops, and the joint publication of results across research projects and in cooperation between senior researchers and doctoral students.

In the first part of this book, we present the overall research goals and questions, and the organizational structure of FORSEC. In the second part, we illustrate the three research clusters of FORSEC, namely PreSTA, STAR and CLOUD, in more detail. In the third and most comprehensive part of this book, we provide a description of all eleven research projects, including their publications in terms of abstract, citation and URL where the full article can be retrieved. In the final reference section, we list all publications of FORSEC in alphabetical order of the first author.

We hope that this report and the set of more than 100 FORSEC publications will stimulate further research on IT security, which we believe will remain one of the most challenging areas in future research on information and communication technologies.

We would like to thank Eva Weishäupl and Dr. Christian Richthammer for their great editorial support.

Regensburg, December 2017 G. Pernul, G. Schryen, R. Schillinger

Coordinators, Principal Investigators and Managing Director

	Prof. Dr. Claudia Eckert TU München Fraunhofer AISEC	Security Architectures for Mobile Equipment (TP1)
	Prof. Dr. Felix Freiling FAU Erlangen-Nürnberg	Software Protection and Anti Forensics (TP5) Security Awareness (TP6)
	Prof. Dr. Hermann de Meer University of Passau	Secure Migration of Virtual Machines (TP4)
	Prof. Dr. Doğan Kesdoğan University of Regensburg	Security and Data Protection in Smart Grids (TP11)
	Prof. Dr. Günther Pernul University of Regensburg	*Coordinator of FORSEC* Identity 3.0 (TP7) Next Generation Online Trust (TP8)
	Prof. Dr. Joachim Posegga University of Passau	IoT Security (TP2) Web Security (TP9)
	Prof. Dr. Hans Peter Reiser University of Passau	Security Concepts for Virtualized Infrastructures (TP3)

Coordinators, Principal Investigators and Managing Director

	Prof. Dr. Guido Schryen University of Regensburg	*Coordinator of FORSEC* Economic Planning and Evaluation of IT Security (TP10)
	Dr. Zinaida Benenson FAU Erlangen-Nürnberg	Security Awareness (TP6)
	Dr. Ing. Tilo Müller FAU Erlangen-Nürnberg	Software Protection and Anti Forensics (TP5)
	Prof. Dr. Rolf Schillinger FH Würzburg-Schweinfurt	*Managing Director of FORSEC*

1. Structure of Alliance and Research

1.1 Introduction

Current phenomena, such as outsourcing, service oriented architectures, cloud computing and also the broad pervasion of every kind of work process with information technology gave rise to a situation that was inconceivable a few years ago. At the turn of the millennium, central IT systems, for example information systems in companies or control systems of public infrastructures (energy grid, traffic control systems), were either isolated or connected to other IT systems within tight, accurately defined boundaries. In the last few years, these boundaries have become increasingly transparent or even disappeared completely. Modern IT systems became versatile, flexible, and highly interconnected, yet fragile constructs. Only a few years ago, anti-virus software and firewalls were considered to provide sufficient protection against attacks on IT systems, but more recently it has become apparent that these measures are obsolete. Complex and adaptive attacks on IT systems (for example Flame, Regin, Pushdo or Gameover ZeuS) demonstrated the capability of abuse and industrial espionage and exposed the weakness of current defensive countermeasures.

Currently, a typical IT security process is a cycle that consists of three phases. Preparedness (P1) describes methods for securing IT systems as well as raising people's awareness of the secure use of IT systems. Research in this area is particularly focused on intrusion prevention and the reduction of the expected damage. Response (P2) is carried out during periods of elevated threat levels or ongoing attacks and includes effective recognition of and subsequent defense from attacks as well as methods for gathering information to support the ensuing conviction of the attackers. Finally, the recovery, auditing and forensics (P3) phase is mostly concerned with methods for the identification of offenders and the recovery of systems and data after successful attacks.

Considering the current threat situation, it becomes apparent that the usually strict separation of the three phases is ineffective. In particular, such an approach does not account for synergies arising from the tight interconnection of phases. Thus, in contrast to viewing all phases separately, FORSEC examines the connections and interfaces between them. In Figure 1, each arrowhead represents the location of one of the examined interfaces. Relations between the phases themselves appear unidirectional due to their cyclic nature. Nonetheless, the information flow along each of these interfaces is bidirectional.

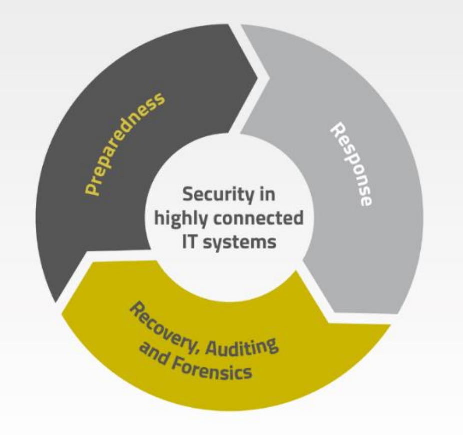
Preparedness
Response
Security in highly connected IT systems
Recovery, Auditing and Forensics

Figure 1: FORSEC IT security process

While it is obvious that preparedness measures are the foundation on which response builds upon, FORSEC additionally considers the possibly important, but yet largely unexplored information flow in the opposite direction, from the response back to preparedness measures. Situational knowledge gained in the recovery phase offers a very important but currently untapped data source for the preparedness phase. Feedback from recovery to response about quantity, quality, and possible additional parameters of the collected data is the last synergy that FORSEC examines. The main expected output of FORSEC, a new, integrated and interdisciplinary concept in the form of an integrated security process for highly connected IT systems, is the direct consequence of these considerations.

In joint research effort, eight working groups from five Bavarian research institutions are involved in the Bavarian research alliance FORSEC: four universities with faculties and departments of different scope (Faculty of Business, Economics, Management Information Systems at the University of Regensburg, Faculty of Computer Science and Mathematics at the University of Passau, Faculty of Computer Science at TU Munich, Faculty of Engineering at FAU Erlangen-Nürnberg), and – associated – the Institute of Applied and Integrated Security (AISEC) at the Fraunhofer Institute in Garching, Munich. An overview of the participating principal investigators along with the research institutions they are associated with and the projects coordinated by them is depicted in Table 2.

1.2 Clusters and Research Questions

FORSEC organizes its research efforts along three specific research questions, each dedicated to one of the interfaces between two phases of the cyclic security process as described in the previous section.

Connecting phases P1 and P3, research question Q1 (cf. Table 1) aims at identifying protection targets, protection goals and associated risks through knowledge on practical weaknesses, security incidents, and associated damages. Such information can be – and quite frequently already is – gathered in the recovery phase. After finishing recovery for a specific security incident, however, the wealth of information is currently only used in aggregated reports and then archived in IT security management systems. In the integrated FORSEC security process, this information is preprocessed and serves as input to the preparedness measures in P1.

During P1, suitable protection mechanisms for the previously identified protection targets, goals and risks are established and implemented as part of standard IT operations procedures. When these defense actions are carried out during response in P2, it is possible that these targets, goals and risks turn out to be incomplete, contradicting or generally unmanageable. Research question Q2 focuses on redefining the interface between P1 and P2 in order to add a feedback channel that provides information on the (un-)suitability of said targets, goals and risks back to P1.

In order to assess security incidents in terms of not only the damage they have caused but also technical efficiency and economic benefit of the implemented protective measures, there is a general consensus on the need for thorough forensic analysis, auditing and controlling within the recovery phase P3. The underlying data for these assessments needs to be captured during P2 but, especially in the case of highly-connected systems, the composition of the data has yet to be investigated within research question Q3.

The final research question Q4 merges the research results of Q1-Q3 into the central, integrated security process for highly connected IT systems.

Table 1: Research questions

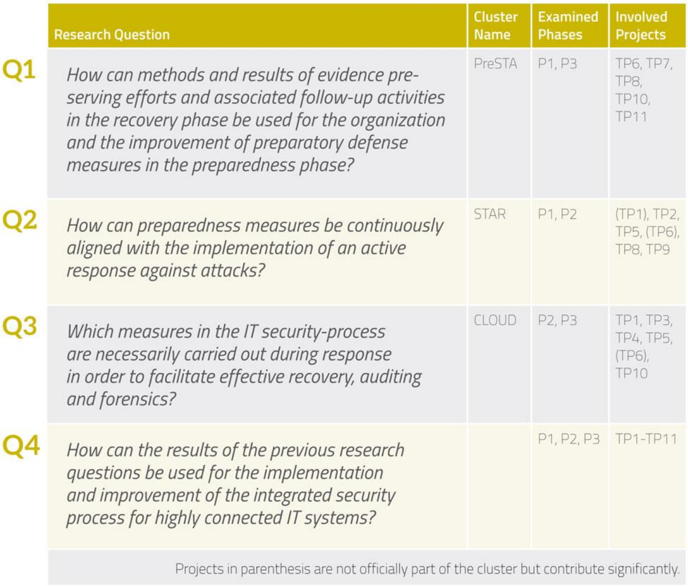

	Research Question	Cluster Name	Examined Phases	Involved Projects
Q1	*How can methods and results of evidence preserving efforts and associated follow-up activities in the recovery phase be used for the organization and the improvement of preparatory defense measures in the preparedness phase?*	PreSTA	P1, P3	TP6, TP7, TP8, TP10, TP11
Q2	*How can preparedness measures be continuously aligned with the implementation of an active response against attacks?*	STAR	P1, P2	(TP1), TP2, TP5, (TP6), TP8, TP9
Q3	*Which measures in the IT security-process are necessarily carried out during response in order to facilitate effective recovery, auditing and forensics?*	CLOUD	P2, P3	TP1, TP3, TP4, TP5, (TP6), TP10
Q4	*How can the results of the previous research questions be used for the implementation and improvement of the integrated security process for highly connected IT systems?*		P1, P2, P3	TP1-TP11

Projects in parenthesis are not officially part of the cluster but contribute significantly.

1.3 Integration of Projects into Clusters and Establishment of Inter-Cluster Connections

FORSEC's research activities have not been commenced from scratch. Instead, the already existing research infrastructure of FORSEC's members is used as a foundation on which the FORSEC alliance is built upon. Situated on this level are eleven projects, each of which covers the central research topics and profits from the expertise of a particular partner. Taking advantage of the benefits a research alliance provides in terms of interdisciplinarity, efficiency and shared expertise over independent or only very loosely integrated projects, the FORSEC alliance was initiated to answer the four research questions introduced in the previous section more efficiently, quicker and – through interdisciplinary collaboration – more thoroughly than would have been possible without FORSEC.

In order to integrate eleven existing projects in a way which allows to address the aforementioned research questions, a bottom-up approach was employed to group the projects into three research clusters. Each of these three research clusters aims at answering one of the research questions Q1-Q3, while all clusters cooperate to answer the fourth, all-encompassing research question Q4. Table 1 lists research questions, the grouping of projects into the three clusters and the mapping of the clusters to the research questions. The overall alliance character of FORSEC is obtained by integrating the three clusters on a scientific level and additionally by coordinating all activities and the non-scientific functions in a global FORSEC office.

From our experience, we have learned that integration of projects into clusters is an excellent measure to quickly initiate and maintain collaboration between projects. In a research alliance, however, collaboration also between clusters is important. Such inter-cluster collaboration is established in FORSEC at three levels. First, the collaborative working environment within FORSEC fosters loose collaboration between individual projects and between individual researchers. Second, the already existing know-how of some projects turned out to be valuable for projects beyond cluster borders, leading to unexpected but very welcome inter-cluster collaborations. TP6 and TP1 are an example of this degree of collaboration since, with regards to content, they are only a member of a single cluster but in reality contribute significantly to other clusters as well. Therefore, they are listed multiple times in Table 1 but appear only in parenthesis in clusters where they are no official member of. Finally, and most formally, transfer projects have been defined between clusters with the explicit mission of connecting the particular clusters. Figure 2 depicts the clustering of projects and the intra-cluster connections taking into account all forms of inter-cluster collaboration currently established in FORSEC.

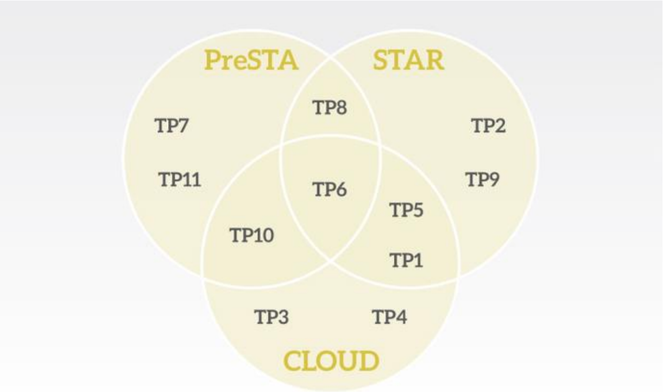

Figure 2: FORSEC clusters and their projects

2. Clusters

As described in Chapter 1 and depicted in Figure 2, FORSEC's research is organized in three research clusters. In this chapter, these three clusters are described in detail. In addition, connections between projects are covered. The description of the individual projects (TP1 – TP11) follows in Chapter 3.

2.1 Cluster I – PreSTA

In accordance with research question Q1, the cluster "Preparatory and Socio Technical Aspects of security in highly connected IT Systems" (PreSTA) examines the organization and improvement of preparedness measures by not only looking at these measures within the boundaries of phase P1 but also studying the numerous links between phases P1 and P3. The five projects which comprise this cluster cover a broad area of research, with technical topics (TP7, TP8, TP11), economical questions (TP10) and human factors questions (TP6) being researched.

Regarding the technical topics, TP7 was concerned with aspects on how to improve the secure handling of users' access in large-scale scenarios. By coming up with structured solutions for several subordinate problems, TP7 contributed to the cluster in various ways. As the second technical TP, TP8 addressed the issue of trust in highly connected IT systems. Most of the currently used trust management systems (e.g. reputation systems) use a rather static approach, leading to high non-transparency and a decreasing level of trust in the system itself. To mitigate these weaknesses, TP8 developed a novel trust management approach based on interactive visualization. Thereby, TP8 benefited from TP6's expertise in user studies and TP7's knowledge in clustering algorithms. Also with a technical focus, TP11 covered the privacy and security relevant aspects of Smart Grid, Smart Meter and IoT appliances inside a household. The solutions of TP11, e.g. regarding visual analysis and usability aspects, were developed in collaboration with TP8 and TP6. TP11 also benefited from TP10's expertise on economic aspects.

As already mentioned, TP10 covered the economic aspects of information security in preparatory defense. Organizations face the challenging task of making decisions with regard to priorities and budgets of investments in security countermeasures by estimating the costs and benefits of possible investments. In cooperation with TP7, TP10 proposed a decision support model to assist decision makers whether to invest in an Identity and Access Management System (IAMS) and if so, in which kind. Regarding human factors questions, TP6 studied which influence users have on IT security, especially with respect to the topics researched by the other TPs. Cooperating with other TPs was a central aspect of TP6.

2.2 Cluster II – STAR

The cluster "Smart & Trusted ARchitectures" (STAR) answers research question Q2, bridging the gap between preparedness and active protection against attacks. The work in Cluster STAR is centered around the smart city prototype with the goal of improving the quality of life of the city's inhabitants: At the lower level, sensors provide data and actuators receive commands. The sensors and actuators are connected to smart platforms which act as gateways to the smart city web services. TP2 built two smart home prototypes to implement this part of the smart city. At the higher level STAR envisions the smart city to provide web services for its inhabitants. These receive sensor data from the smart homes and make them available on an open platform to be aggregated, analyzed, etc. These web services were implemented as part of TP9. A major topic in the smart city prototype was security: The smart home prototypes provide a detailed access control system with different roles and rights; in addition, users can specify using security policies how sensor data provided to the smart city may be used. TP9 implemented mechanisms in the web services for enforcing the security policies attached to sensor data. In a similar direction, TP2 also developed a taint analysis system to track how apps in the smart home use sensitive data. In a smart city, participants may also attempt to provide incorrect sensor data for various reasons; TP8's expertise in recommendation systems helps to identify such attacks, thus adding another layer of security. In contrast to TP2 and TP9, which mainly take the perspective of protecting smart city users, TP5 adds the perspective of protecting developers of apps in the smart city: Their obfuscation and anti-analysis techniques protect intellectual property in apps from being compromised. Thus, the cluster STAR combines secure architectures, i.e. preparedness, with active protection against attacks.

All constituting projects are of a highly technical nature and their contributions to STAR are explained in Chapter 3. TP1 has moved to cluster CLOUD but continues to provide its research results to the TPs in STAR that build upon them (mostly TP2 and TP9).

2.3 Cluster III – CLOUD

The cluster CLOUD is primarily focused on security questions in a cloud environment. In order to answer research question Q3, it examines the complex links between the recovery and response phases, where it is important to minimize the attack potential while at the same time maximizing the possibilities for detailed analyses and thorough evidence preservation. As such, it tackles the challenge of integrating lightweight incident detection methods that can be used permanently on production systems with potentially heavyweight methods that yield a detailed in-depth view of incident and accurate preservation of evidence, taking into account Quality of Service (QoS) requirements, security policies defined by cloud provider and cloud customer, and economical constraints.

The cluster addressed the research question by designing a joint cloud architecture. The CloudIDEA architecture is a practical solution that implements a trade-off between the requirements of being resource efficient and providing forensic means in case of an incident (Fischer et al. 2015; Taubmann, Reiser, et al. 2015). Additionally, the cluster discussed the threats to IaaS based cloud computing where a classification of cloud specific malware attacks was systematically analyzed in order to describe a state-of-the-art attacker model (Rakotondravony, Taubmann, et al. 2017).

While TP3, TP4 and TP5 as well as TP1 (which moved from STAR to CLOUD cluster during the project) are very technical in their approach, TP10 adds above-mentioned economic considerations concerning business continuity and, more broadly, attack cost. Cluster members TP4 and TP10 collaborated on the discussion of the relation between QoS, privacy, and security issues in IaaS clouds and the financial impacts of non-compliance to SLAs (Mandarawi and Weishäupl, 2017; Mandarawi, Fischer, et al., 2015). Additionally, TP1 and TP3 collaborated on a joint architecture to run and monitor malware in cloud environments in order to build and train a resource efficient machine learning based detection mechanism (Taubmann and Kolosnjaj, 2017).

3. Projects

This chapter is concerned with the projects (TPs) that make up the clusters presented in the previous chapter. All project descriptions follow the same logic, starting with the project's overall goal, results achieved, contribution to the FORSEC research alliance, and outlook beyond FORSEC. Finally, the abstracts of all publications of the respective project are provided along with their citations and with URLs provided by the publishers where available.

Table 2 provides an overview of FORSEC's projects and the responsible universities.

Table 2: FORSEC projects and responsible universities

#	Title	Responsible University
TP1	Security Architectures for Mobile Equipment	TU München
TP2	IoT Security	University of Passau
TP3	Security Concepts for Virtualized Infrastructures	University of Passau
TP4	Secure Migration of Virtual Machines	University of Passau
TP5	Software Protection and Anti Forensics	FAU Erlangen-Nürnberg
TP6	Security Awareness	FAU Erlangen-Nürnberg
TP7	Identity 3.0	University of Regensburg
TP8	Next Generation Online Trust	University of Regensburg
TP9	Web Security	University of Passau
TP10	Economic Planning and Evaluation of IT Security	University of Regensburg
TP11	Security and Data Protection in Smart Grids	University of Regensburg

3.1 TP1 – Security Architecture for Mobile Equipment

3.1.1 Project Overview

The goal of our project is to develop approaches for improving applicable for security on mobile architectures. We have evaluated the security aspects for the ARM architecture and its virtualization extensions. In order to keep up with the advanced attacks, such as code reuse attacks, we extend the defense to dynamic binary analysis with virtualization extensions on ARM architecture. Furthermore, we develop and extend machine learning approaches to detect and triage malware in an anomaly detection framework, based on the data gathered from static and dynamic malware analysis. Using nonparametric topic modeling, convolutional and recurrent neural networks we enable fast malware detection, similarity search, classification and actor attribution, in large-scale datasets. In the next iteration, we adapt our anomaly detection approaches for the resource-constrained environment, such as mobile devices.

3.1.2 Results Achieved

Evaluating the ARM architecture as well as its security and virtualization extensions on the new version ARM processors is an interesting topics on the dynamically binary analysis field. Currently, especially with the promotion of advanced attacks methodology, traditionally static binary analysis cannot keep abreast with the advanced attacks, especially code reuse attacks. Therefore, we extend traditional static binary analysis (SBA) method on Intel X86/X64 architecture to dynamical binary analysis(DBA) with virtualization extensions and then to ARM architecture with virtualization. We develop DRAKVUF, a virtualization based agentless black-box binary analysis system (as part) and VMI4ddCRAs, a binary rewriting tool based on DRAKVUF to defend against code reuse attacks.

DRAKVUF is a virtualization based agentless black-box binary analysis system. It allows for in-depth execution tracing of arbitrary binaries (including operating systems), all without having to install any special software within the virtual machine used for analysis. We published this work on the ASCAS-2014 conference as the summarization of our works on virtualization topic as well as the virtualization-based binary analysis technique.

On the ARM platform, libvmi was also updated to be able to understand ARM paging. For Xen we implemented the memory tracing system using SLAT and it was merged into mainline Xen in the 4.6 release. We also described a theoretical security system different from all current approaches on ARM. It did not got fully implemented but that was the goal we are working on.

In addition, in order to extend of VMI4daCRAs to Android ecosystem, we also do some basic works. VMI4daCRAs framework is the binary rewriting based way to defend against the code reuse attacks from the hypervisor. It cannot understand the Android semantic, especially the bytecode (Android APK) and the JIT-generated code. Therefore, currently we design and implement a compiler-based way to mitigate the code reuse attacks on ARM based Android system. We name it as daVTDroid. daVTDroid can defend against nearly all vTable hijacking, not only vTable corruption hijacking attacks and vTable injection attacks, but also the advanced vTable reuse attacks. Although we did not verify original COOP attacks can be available on ARM based Android system, we believe the statement which says COOP would be easily extended to all RISC architectures which use registers as the parameters transformation.

Next, we develop and extend machine learning approaches to detect and triage malware in an anomaly detection framework, based on the data gathered from static and dynamic analysis. We gather large-scale datasets of malware system calls, instruction sequences, PE Header as well as Rich Header features. Based on this data we develop various machine learning approaches for large-scale analysis. Our methodology is directed not only towards a better detection accuracy, but also for robustness and adaptivity, meaning that we can retrain the model efficiently in case of a high influx of malware samples.

We adapt nonparametric topic modeling methodology for semi-supervised learning in order to capture the groups of system calls responsible to particular aspects of malware activity. Topic modelling enables us to make our detection system semantics-aware, meaning that the activities detected using this statistical model are interpretable for a malware analyst. The features obtained using topic modelling are combined with the data from PE header in a unified framework for data-driven malware analysis, published in a conference paper. Furthermore, we develop approaches for a more in depth analysis of malware instruction sequences that identifies code patches common for various malware families. Using a framework based on convolutional networks we improve malware classification, as we are able to extract the discriminative parts of the code location-independently. This work is based on previous papers that use image processing methods to identify interesting patterns in malware code. The approach we use here is also robust on small code obfuscations, such as instruction reordering and adding bogus code.

On the other hand, we dedicate efforts to enable fast nearest neighbor search on a large set of malware samples. Gathering of some types of data, such as dynamic malware analysis results or instruction sequences can be slow and cumbersome, and metadata from the PE header can and very often is stripped or tampered with. However, we found that the data from a special part of the Windows executables called Rich Header is very often left unchanged by malware authors, and it carries very useful information about the compilation settings used when building the malicious executables. This header can be decoded very fast on a large number of samples, which enables us to quickly preprocess data prior to similarity search. Based on stacked autoencoders and the ball tree data structure we create a system where we identify nearest neighbour malware samples on various granularity levels and identify actors in a matter of milliseconds on a million-scale dataset. This in turn enabled us to detect members of some contemporary malware families that previously only had generic antivirus signatures.

In the next stage, we adapt our anomaly detection approaches with compressed and lightweight models, enabling us to detect malicious activity in a resource-constrained environment, such as mobile or Internet of Things (IoT) devices. There are many types of constraints that these devices impose. We consider three broad types of anomaly detection scenarios: communication efficiency, model size efficiency and data acquisition efficiency.

We have developed a communication-efficient learning procedure applicable to a client-server scenario, for example intrusion detection based on real-time data from multiple mobile or IoT devices, or other kind of collaborative anomaly detection. Our framework enables reliable identification of malicious and inaccurate clients from the labelled data they provide. This enables the selection of a minimal set of clients that provides sufficient learning performance in terms of classification accuracy. Using a combination of standard machine learning models, such as logistic regression, and sparsity constraints on the weight of different labelled data sources, we can achieve results close to the results when using all clients, both on real and synthetic datasets, while using data from less than a half of clients.

Furthermore, we adapt online learning approaches to a scenario of low memory budget. This means that, in case of nonparametric machine learning methods where model grows with the scale of gathered data, we control this growth using an optimization procedure, thereby limiting the memory footprint. This strategy is called budgeted learning and has been previously applied in image processing. We test this approach on a task of continuous authentication on mobile devices. The experiments that we executed on 28 subjects show that we can limit the size of the authentication model, for instance by controlling the support vectors in a kernel-based system. This kind of model efficiently determines if the behaviour in terms of touchscreen interaction and sensor measurements belongs to the valid user of the mobile device. However, this methodology can also be used in other scenarios, such as intrusion detection on memory-constrained devices.

The last resource-constrained scenario is related to the constraint in data acquisition. In many scenarios related to security, we monitor computer systems by tracing the behaviour of programs or users. In particular, in case of Virtual Machine Introspection, it is very resource-intensive to trace and record all the system calls and network traffic in a cloud environment. Motivated by this constraint, we evaluate approaches for the optimal selection of the subset in the set of traceable events that would enable sufficient malware detection performance, while maintaining the low tracing overhead. We choose to test a neural network with attention model, where attention vector determines which part of the input features are important for modelling the malware execution sequences. By imposing sparsity on the attention vector we minimize the input features needed to train and test our models.

3.1.3 Contribution to FORSEC Research Alliance

Our project belongs to the cluster 3 called CLOUD, dedicated to improving security in the area of cloud computing systems. Our work related to Virtual Machine Introspection has a strong relation to this scenario. VMI techniques are an essential tool in monitoring virtual machines in the cloud environment, as they enable acquisition of data useful for intrusion detection or forensics. Furthermore, the machine learning approaches that we developed are very useful in combination with these monitoring techniques for malware detection and analysis in the cloud environment. We have contributed to two papers that were created as a collaboration of all the cluster members.

The first paper named "CloudIdea" represents an architecture that encompasses all of the security measures needed in the cloud and includes the contribution of all the cluster members in terms of methods and techniques useful in this scenario. The paper includes instructions for integrating lightweight intrusion detection methods that can continuously be used in the cloud computing systems with heavyweight methods that provide more detailed information about incidents, but are more resource-intensive, while taking into account Quality of Service (QoS) requirements, security policies defined by cloud provider and cloud customer, and economical constraints. We have, in particular, contributed by describing the state of the art in Virtual Machine Introspection methods for monitoring and malware analysis, as well as in introducing the role of machine learning and other statistical methods in processing the data that we can get from the virtual machine monitoring. The second paper, currently under review, represents a survey of different threats in the cloud environment, describing different types of malware and attacks that can endanger the proper functioning of IaaS systems. During our FORSEC project we maintained the closest collaboration with TP3, as both of our groups are interested in developing systems relying on the VMI concept. Therefore we had multiple discussion sessions about this topic, resulting in an idea for a joint research paper. We have written a paper draft, currently under review, describing our approach of Virtual Machine Introspection that takes into account resource constraints imposed by tracing a large set of events. In our approach we use Machine Learning to adapt the set of events that we need to trace in order to maintain optimal intrusion detection performance, while keeping the overhead minimal.

3.1.4 Beyond FORSEC

VMI4ddCRAs is a framework to defend against code reuse attacks with binary rewriting technique. Due to the product of dex2oat tool, nearly all apps on Android system (bytecode form) should be transformed to OAT file in Android ART runtime. OAT is a dynamic (shared) library in Android and it can be rewritten by VMI4ddCRAs. But for the apps which have JIT features, VMI4ddCRAs cannot rewrite these JIT-generated code and of course cannot defend against JIT-ROP and JIT-COOP attacks. Therefore, we want to extend VMI4ddCRAs in order to support the whole Android OS. At the final step of our VMI4ddCRAs framework, we want to use it as a whole framework to mitigate code reuse attacks with the form of binary rewriting.

On the side of machine learning, we are planning to further adapt our work in large-scale malware detection and analysis to Android malware, in order to further improve the security of mobile devices. Currently we are working on the method for fast matching of control flow graphs and function call graphs of Android apps, which would enable us to do fast triage of Android malware. Furthermore, we are exploring the effects of adversarial attacks on our machine learning systems. This kind of systems are vulnerable to both exploratory and causative attacks, meaning that small changes in the training or test set can endanger their performance. We need to explore the game-theoretic scenario of optimal attack and defense in adversarial learning in case of malware detection systems.

3.1.5 Publications

Kolosnjaji and Eckert (2015a): Neural Network-Based User-Independent Physical Activity Recognition for Mobile Devices	
Abstract	Activity recognition using sensors of mobile devices is a topic of interest of many research efforts. It has been established that user-specific training gives good accuracy in accelerometer-based activity recognition. In this paper we test a different approach: offline user-independent activity recognition based on pretrained neural networks with Dropout. Apart from satisfactory recognition accuracy that we prove in our tests, we foresee possible advantages in removing the need for users to provide labeled data and also in the security of the system. These advantages can be the reason for applying this approach in practice, not only in mobile phones but also in other embedded devices.
Citation	Kolosnjaji, B., and Eckert, C. 2015a. "Neural Network-Based User-Independent Physical Activity Recognition for Mobile Devices," in *Proceedings of the 16th Conference on Intelligent Data Engineering and Automated Learning (IDEAL 2015)*, Wroclaw, Poland: Springer, pp. 378–386.
URL	https://link.springer.com/chapter/10.1007/978-3-319-24834-9_44

Kolosnjaji and Eckert (2015b): Leveraging Deep Learning for Malware Detection and Classification	
Abstract	As signature-based malware detection systems are unable to cope with the increasing number and variety of malware samples, machine learning has been proposed as a robust alternative. Neural networks have been used in numerous research efforts as a machine learning-based method for the detection and classification of malware, for the purpose of both network-based and host-based intrusion detection. The most used configuration of neural network in these efforts was a perceptron with one hidden layer. However, recent years have brought a significant advancement in neural networks, with new training methods and improved configuration possibilities for neural network units. The advancement is centered around the paradigm of deep learning. This paper contains a description of these new approaches and discusses the possibilities of their application to malware detection and classification problems. A novel malware detection architecture is presented that leverages these advancements for classifying malware based on inputs from static and dynamic analysis results.
Citation	Kolosnjaji, B., and Eckert, C. 2015b. "Leveraging Deep Learning for Malware Detection and Classification," in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015)*, Vienna, Austria.

Kolosnjaji et al. (2016): Adaptive Semantics-Aware Malware Classification	
Abstract	Automatic malware classification is an essential improvement over the widely-deployed detection procedures using manual signatures or heuristics. Although there exists an abundance of methods for collecting static and behavioral malware data, there is a lack of adequate tools for analysis based on these collected features. Machine learning is a statistical solution to the automatic classification of malware variants based on heterogeneous information gathered by investigating malware code and behavioral traces. However, the recent increase in variety of malware instances requires further development of effective and scalable automation for malware classification and analysis processes. In this paper, we investigate the topic modeling approaches as semantics-aware solutions to the classification of malware based on logs from dynamic malware analysis. We combine results of static and dynamic analysis to increase the reliability of inferred class labels. We utilize a semi-supervised learning architecture to make use of unlabeled data in classification. Using a nonparametric machine learning approach to topic modeling we design and implement a scalable solution while maintaining advantages of semantics-aware analysis. The outcomes of our experiments reveal that our approach brings a new and improved solution to the reoccurring problems in malware classification and analysis.
Citation	Kolosnjaji, B., Zarras, A., Lengyel, T., Webster, G., and Eckert, C. 2016. “Adaptive Semantics-Aware Malware Classification,” in *Proceedings of the 13th Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA 2016)*, San Sebastian, Spain.
URL	https://link.springer.com/chapter/10.1007/978-3-319-40667-1_21

Kolosnjaji et al. (2017): Empowering Convolutional Networks for Malware Classification and Analysis	
Abstract	Performing large-scale malware classification is increasingly becoming a critical step in malware analytics as the number and variety of malware samples is rapidly growing. Statistical machine learning constitutes an appealing method to cope with this increase as it can use mathematical tools to extract information out of large-scale datasets and produce interpretable models. This has motivated a surge of scientific work in developing machine learning methods for detection and classification of malicious executables. However, an optimal method for extracting the most informative features for different malware families, with the final goal of malware classification, is yet to be found. Fortunately, neural networks have evolved to the state that they can surpass the limitations of other methods in terms of hierarchical feature extraction. Consequently, neural networks can now offer superior classification accuracy in many domains such as computer vision and natural language processing. In this paper, we transfer the performance improvements achieved in the area of neural networks to model the execution sequences of disassembled malicious binaries. We implement a neural network that consists of convolutional and feedforward neural constructs. This architecture embodies a hierarchical feature extraction approach that combines convolution of n-grams of instructions with plain vectorization of features derived from the headers of the Portable Executable (PE) files. Our evaluation results demonstrate that our approach outperforms baseline methods, such as simple Feedforward Neural Networks and Support Vector Machines, as we achieve 93% on precision and recall, even in case of obfuscations in the data.
Citation	Kolosnjaji, B., Eraisha, G., Webster, G., Zarras, A., and Eckert, C. 2017. "Empowering Convolutional Networks for Malware Classification and Analysis," in *Proceedings of the 30th International Joint Conference on Neural Networks (IJCNN 2017),* Anchorage, AK.
URL	http://ieeexplore.ieee.org/document/7966340/

Lengyel, Kittel, et al. (2014): Pitfalls of Virtual Machine Introspection on Modern Hardware	
Abstract	Over the last few years there has been immense progress in developing powerful security tools based on Virtual Machine Introspection (VMI). VMI offers unique capabilities which can be used to check and enforce security policies in the presence of a potentially compromised guest. With the introduction of new hardware virtualization extensions, VMI can be further enhanced to provide lightweight, in-band control over the execution of virtual machines. In publications released before the extensions were available, security researchers issued warnings that these new extensions may be used to subvert VMI. Since hardware supporting these extensions is now available, in this paper, we aim to discuss and re-evaluate claims made in prior-art. We further continue the discussion by highlighting critical limitations of the virtualization extensions. We go on to show that thorough consideration and understanding of these limitations is necessary when developing VMI based security applications. Otherwise, improper handling will inadvertently expose these applications to subversion attacks. Finally, we take a look at Intel's normal and dual-monitor System Management Mode and discuss how they can be used to both implement and subvert VMI based security applications.
Citation	Lengyel, T., Kittel, T., Webster, G., and Torrey, J. 2014. “Pitfalls of Virtual Machine Introspection on Modern Hardware,” in *Proceedings of the 1st Workshop on Malware Memory Forensics (MMF 2014)*, New Orleans, LA, December (available at https://www.sec.in.tum.de/assets/Uploads/pitfalls-virtual-machine.pdf).

Lengyel, Maresca, et al. (2014): Scalability, Fidelity and Stealth in the DRAKVUF Dynamic Malware Analysis System	
Abstract	Malware is one of the biggest security threats on the Internet today and deploying effective defensive solutions requires the rapid analysis of a continuously increasing number of malware samples. With the proliferation of metamorphic malware the analysis is further complicated as the efficacy of signature-based static analysis systems is greatly reduced. While dynamic malware analysis is an effective alternative, the approach faces significant challenges as the ever increasing number of samples requiring analysis places a burden on hardware resources. At the same time modern malware can both detect the monitoring environment and hide in unmonitored corners of the system. In this paper we present DRAKVUF, a novel dynamic malware analysis system designed to address these challenges by building on the latest hardware virtualization extensions and the Xen hypervisor. We present a technique for improving stealth by initiating the execution of malware samples without leaving any trace in the analysis machine. We also present novel techniques to eliminate blind-spots created by kernel-mode rootkits by extending the scope of monitoring to include kernel internal functions, and to monitor file-system accesses through the kernel's heap allocations. With extensive tests performed on recent malware samples we show that DRAKVUF achieves significant improvements in conserving hardware resources while providing a stealthy, in-depth view into the behavior of modern malware.
Citation	Lengyel, T., Maresca, S., Payne, B. D., Webster, G. D., Vogl, S., and Kiayias, A. 2014. “Scalability, Fidelity and Stealth in the DRAKVUF Dynamic Malware Analysis System,” in *Proceedings of the 30th Annual Computer Security Applications Conference (ACSAC2014),* New Orleans, LA.
URL	https://dl.acm.org/citation.cfm?id=2664252

Lengyel et al. (2015): Virtual Machine Introspection with Xen on ARM	
Abstract	In the recent years, virtual machine introspection has become a valuable technique for developing security applications for virtualized environments. With the increasing popularity of the ARM architecture and the recent addition of hardware virtualization extensions there is a growing need for porting existing tools to this new platform. Porting these applications requires proper hypervisor support, which we have been exploring and developing for the upcoming Xen 4.6 release. In this paper we explore using ARM's two-stage paging mechanisms with Xen to enable stealthy, efficient tracing of guest operating systems for security purposes.
Citation	Lengyel, T., Kittel, T., and Eckert, C. 2015. “Virtual Machine Introspection with Xen on ARM,” in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015)*, Vienna, Austria.

Webster et al. (2017): Finding the Needle. A Study of the PE32 Rich Header and Respective Malware Triage	
Abstract	Performing triage of malicious samples is a critical step in security analysis and mitigation development. Unfortunately, the obfuscation and outright removal of information contained in samples makes this a monumentally challenging task. However, the widely used Portable Executable file format (PE32), a data structure used by the Windows OS to handle executable code, contains hidden information that can provide a security analyst with an upper hand. In this paper, we perform the first accurate assessment of the hidden PE32 field known as the Rich Header and describe how to extract the data that it clandestinely contains. We study 964,816 malware samples and demonstrate how the information contained in the Rich Header can be leveraged to perform rapid triage across millions of samples, including packed and obfuscated binaries. We first show how to quickly identify post-modified and obfuscated binaries through anomalies in the header. Next, we exhibit the Rich Header's utility in triage by presenting a proof of concept similarity matching algorithm which is solely based on the contents of the Rich Header. With our algorithm we demonstrate how the contents of the Rich Header can be used to identify similar malware, different versions of malware, and when malware has been built under different build environment; revealing potentially distinct actors. Furthermore, we are able to perform these operations in near real-time, less than 6.73 ms on commodity hardware across our studied samples. In conclusion, we establish that this little-studied header in the PE32 format is a valuable asset for security analysts and has a breadth of future potential.
Citation	Webster, G. D., Kolosnjaji, B., Pentz, C. von, Kirsch, J., Hanif, Z. D., Zarras, A., and Eckert, C. 2017. “Finding the Needle: A Study of the PE32 Rich Header and Respective Malware Triage,” in *Proceedings of the 14th Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA 2017)*, Bonn, Germany.
URL	https://link.springer.com/chapter/10.1007/978-3-319-60876-1_6

3.2 TP2 – Internet of Things Security

3.2.1 Project Overview

TP2 Internet of Things Security explores new security measures to make the Internet of Things more secure in a smart home and smart city scenario. Therefore, TP2 has two main focuses: (1) specification and implementation of a prototype of a smart home and smart city, and (2) research on new security technologies beyond the state of the art using this prototype. At the lower level this includes authentication mechanisms between sensors and embedded platforms, as well as between platforms and smartphones. We also evaluated the feasibility of generating digital signatures on sensors. At the platform level, Flowcoaster extended existing work on taint analysis.

3.2.2 Results Achieved

Secure Smart Home

The results of TP2 are centered around the smart home which also lies at the core of the Cluster STAR scenario. The smart home connects the following components: smart platforms, sensors and actuators, and smartphones. The smart platforms are built on Odroid XU4s running Android; this allows for easy interoperability with the smart home app on Android smartphones. They form the core of our smart home. The app on the smartphone acts as a user interface to interact with the smart platform (we assume that every smart home inhabitant has a smartphone). The sensors and actuators deliver sensed data to the smart platform and act on commands received respectively. Currently our smart home supports connecting to a sensor or actuator in a variety of ways, e.g. locally via GPIO pins, locally via a serial USB connection, remotely via the (wireless) network connection.

The smart platform can also be connected to our smart city platform SERIOS which is explained in more detail in TP9. It then pushes sensor data to SERIOS as specified by the user.

The code for this secure smart home is publicly available (Werli et al. 2016).

Authentication of smart home devices with QR codes (Marktscheffel et al. 2016)

The first challenge to building a secure smart home is the setup process: How can the components, i.e. the Odroids and smartphones, be connected and authenticated to each other to allow encrypted communication? Furthermore, in a smart home, we cannot assume that there are only expert users; the smart home functionality shall be usable for all inhabitants. Therefore, easy usability was also a requirement for our authentication process.

Our solution is based on QR codes. Thus, we need one device with a screen to display the QR code, and another device with a camera to record the QR code. In terms of authentication, we have the master device managing the smart home and the new device, which is to be registered. We designed the authentication protocol in a flexible way so that either device, master or new device, can be the device recording the QR code or displaying it. While with smartphones we can generally assume that they have a camera and a display, this is not the case with Odroids. In our standard setup only one Odroid is connected to a camera; however, if necessary, the Odroid can also manually be connected to a keyboard and screen. Nevertheless, the flexibility in the protocol helps to reduce the number of situations where the Odroid needs to be connected to a screen.

On a more technical level, the protocol we designed works as follows: When an authorized user requests to connect a new device, a QR code containing the following information is generated: token (array of random bytes), address of the device generating the QR code, and a fingerprint (cryptographic hash of public key). The device scanning the QR code then extracts address and requests the certificate of this device. It then verifies that the fingerprint encoded in the QR code belongs to this certificate. Next, the device with the camera generates a nonce to be used as a challenge. Using this public/private key pair, the token from the QR code, and the nonce, the devices then mutually authenticate. Authentication of the device with the screen is based on the knowledge of information contained only in the QR code (the user has to make sure that only this device can scan the QR code). The device with the screen is authenticated based on proving that it has the private key for the certificate fingerprinted in the QR code. As a QR code for authentication is only created upon request of a user, and each QR code is only used once, replay attacks will not be successful against our protocol. The two channel approach with network communication and the optical channel via the QR code also rules out MITM attacks. Thus, under the assumption of an active network-based attacker, this protocol is secure. However, the attacker must not compromise the devices involved or access the QR code.

Authentication of sensors with minimal user interface

While the protocol above is well suited to authenticate embedded devices and smartphones, this is not the case for sensors and authentication between sensor and Odroid. Sensor and actuators generally do not have screens or cameras. In their most basic forms, e.g. a door buzzer connected to a GPIO pin, they have no user interface at all, just the capability to receive commands (actuator) or send sensed values (sensor). With no user interface, it is not possible to achieve authentication. We explored what the minimal set of user interface components is to support authentication. We are aware of related work based on smartcards; however, they are often too expensive to be included in sensors for smart homes and rely on centralized infrastructure, and therefore we wanted to develop an alternative solution for smart homes.

We show that the following components on the sensor are sufficient for mutually authenticating a sensor and an embedded platform: One button (or similar input device) and one LED. In addition, we require the user to act as a second data channel between sensor and platform. To start authentication, the user has to press the button on the sensor to put the sensor into setup mode. This ensures that once the sensor is connected, it keeps this connection; an attacker must not have physical access to the sensor. Next, the sensor initiates a Diffie-Hellman key exchange with the platform. This step is still vulnerable to Man-in-the-middle or impersonation attacks. Therefore, in the next step, sensor and platform verify with the help of the user that no attacker interfered in the key exchange: The embedded platform generates a random pattern of pressing and releasing the button over a specific time. The user is then requested to press and release the button in this way, which the sensor records. The sensor then sends the recorded pattern (encrypted) to the platform, which verifies if it is correct. If the platform does not have a display, a smartphone securely connected to it can be used alternatively. If the sensor has a display, the verification can be abbreviated by letting the user compare fingerprints of the keys on sensor and platform.

Signatures for sensor data

Building on these authentication mechanisms, we designed methods to support redactable signatures on sensors (Frädrich et al. 2016). Signatures ensure integrity and authenticity of data on its way to the (smart city) cloud; redactable signatures additionally increase privacy by allowing to redact parts of the sensor data. This is useful in cases when not all recipients are authorized to see all parts of the sensor data. Existing algorithms to create redactable signatures required many checks for a hash value being prime, i.e. more computational power than our sensors had. With a change from the standard model to the random oracle model, we now only need to check if a number is odd, with the consequence that our modified algorithm was 1000 times faster than the original. Nevertheless, runtime to sign a 4-part message was still over half a day on a sensor (Zolertia ReMote). We also evaluated different algorithms for digital signatures based on elliptic curve cryptography on sensors (Bauer et al. 2016). Here energy consumption or runtime overhead is not the limiting factor, but RAM size is. Also, hardware acceleration only results in a moderate speedup. On average, the fastest algorithm took 537 and 595ms for signing and verifying respectively. Thus, we conclude that it is possible to generate signatures and even redactable signatures on sensors, but to achieve an efficient solution requires paying attention to the details during implementation.

Taint Analysis for sensitive sensor data on Android
Android smartphones handle lots of sensitive data, e.g. location or contact data. This is even more the case for our smart home platforms: They collect various sensor data which contains private information about the home's inhabitants. Taint analysis, i.e. tracking how an app uses data, is an approach handle this problem: With taint analysis, the platform can enforce that an app uses certain data items only locally, but does not leak it outside of the platform. Taintdroid, a well known approach for taint analysis on Android, lacks support for native code, which can be part of apps.

Therefore, we extended Taintdroid in our approach called Flowcoaster to support tracking of native code. Taintdroid can be efficiently implemented by modifying the Dalvik virtual machine. However, native code is executed directly by the CPU which cannot be modified (in the scope of FORSEC). In Flowcoaster, we use Valgrind; more precisely we modified Taintgrind, a Valgrind tool, for the purpose of our taint analysis. Valgrind uses an intermediate representation (IR) for instrumentation of code; before and after instrumentation code has to be translated to IR and back to binary code for execution. This translation process and Valgrind's memory management, even without any active instrumentation, cause execution to be slower by a factor of 4. Furthermore, Valgrind's architecture implies also that it must have full control over the whole process, i.e. it is not possible in Flowcoaster to switch on instrumentation only when native code is executed. Therefore, Flowcoaster separates execution into two different processes: During app startup, a separate process for handling native code is started which is called Wrapper. The Wrapper includes a Binder service, so that the Dalvik VM can use Binder inter-process-communication to send requests to the Wrapper. This inter-process-communication is required for every switch between Dalvik VM and native code, i.e. whenever a native function is called or returns, and whenever native code calls a function in the JNI interface or it returns. The Dalvik VM also needs to notify the Wrapper when a new shared library or function needs to be loaded. In addition, Flowcoaster needs to keep track of different threads inside the Dalvik VM, so that the Wrapper can correctly match execution requests and results to the correct thread. While at the moment, Flowcoaster is implemented to extend Taintdroid using the Dalvik VM, its flexible architecture would also compatible with other tools or the newer Android runtime. With Flowcoaster, we can show that taint analysis can be extended to cover native code, but only at the price of being much less efficient than for byte code only.

Native Code Security: Security Analysis in untrusted runtime environment

In the field of native security in Android, we consider the following research question: Given that native code has full access to all process memory, how does that affect security tools based on instrumenting the ART runtime or Dalvik VM respectively? Can malicious code attack such analysis tools using native code? This new work is a result of our work on Flowcoaster, where we chose to avoid this potential vulnerability by creating the separate Wrapper process.

Our initial work on this topic shows that with the help of the Android runtime, we can find and modify many internal data structures of the ART runtime. For Taintdroid, we can also show how native code can modify the internal data structures of Taintdroid maliciously. Unfortunately, for most other security analysis tools, their authors did not publish their source code, so that we are not able to assess how easily they can be attacked.

Sensolatr – Simulating Sensors for the Smart City

We currently only have two physical smart homes; this is not sufficient for a smart city on larger scale. This gap is filled by Sensolatr: Sensolatr is a set of scripts to simulate sensor data and push it to SEDARI. It supports a number of different profiles for different scenarios. A recorded history can be played back, optionally with some variations. It can send constant or random values, indicating a faulty sensor. The profiles are stored are in JSON format; thus, it is easy to create new profiles and for Sensolatr to process them. Sensolatr also supports simulation of the library occupation sensors (cf. TP9).

3.2.3 Contribution to FORSEC Research Alliance

The common scenario of cluster STAR is the smart city: Smart homes and other entities with sensors provide data and act upon it; the web platforms of the smart city collect and make available this data to improve the inhabitants' quality of life. Thus, TP2 Internet of Things Security is very much at the core of this cluster. The smart home prototypes built in TP2 are the lower layer of the smart city. As the server-side components of the smart city were developed by TP9, we collaborated closely on the connection of the smart home to the smart city platform SEDARI. Also involved in the development of SEDARI, in particular in the identity management component, was the H2020 European project COMPOSE. The service platform in the smart homes is able to push sensor data to the smart city. The sensor data is encoded in a JSON format. The values of a sensor over time are grouped into a stream, and a service object can group streams of related sensors together, e.g. temperatures for kitchen and living room. The smart home's rights management defines which sensor data must stay in the smart home and which may be sent to the smart city; it can also attach security policies to the data. SEDARI will then enforce these policies. Thus, we have built a prototype covering all layers from sensors, via (local) smart platforms to web-based platforms. This collaboration was very valuable as it allows us to show the benefits of integrating security over several layers.

We also collaborated with TP9 and ACTLab at the University of Passau on the SmartLibrary project: Also for the sensors of the SmartLibrary, Sensolatr is able to simulate the occupation of additional seats in the library. This simulation helps in evaluating the approach to increase the occupany rate of seats in the library.

Our work on implementing redactable signatures (Frädrich et al. 2016) was done in collaboration with TP3. Assurance of data origin and integrity is not only important in sensors, but also as the data is passed on to IoT platforms and cloud nodes.

3.2.4 Beyond FORSEC

The smart city prototype was built in collaboration with TP9 and the H2020 project AGILE. It has been very helpful to show technologies integrated over several layers improve security. Therefore, we plan to continue using this prototype for further projects and as a demonstration platform for upcoming research.

Furthermore, our work on Flowcoaster raised new interesting research questions such as: Can native code undermine the integrity of runtime-based security tools analyzing an app? We started to work on this topic (see above on native security), and plan an initial publication towards the end of FORSEC.

More generally, we plan to continue our research on the architecture and design of secure platforms for the internet of things. Our work on Android security, and in particular Flowcoaster and native code security, showed that generally Android is very well suited to support runtime-based security tools, but there are also some major pitfalls, e.g. native code. Therefore, we plan to research how such pitfalls can be avoided in future platforms, so that they support security tools better.

We will also continue our research on authentication in the Internet of Things. We are currently working on an alternative approach for authentication based on using visual light transmission as a second channel, instead of the user having to press a button. In addition, we plan to expand on our research on authentication of sensors in future research project.

3.2.5 Publications

Bauer et al. (2016): ECDSA on Things: IoT Integrity Protection in Practise	
Abstract	This paper documents some experiences and lessons learned during the development of an IoT security application for the EU-funded project RERUM. The application provides sensor data with end-to-end integrity protection through elliptic curve digital signatures (ECDSA). Here, our focus is on the cost in terms of hardware, runtime and power-consumption in a real-world trials scenario. We show that providing signed sensor data has little impact on the overall power consumption. We present the experiences that we made with different ECDSA implementations. Hardware accelerated signing can further reduce the costs in terms of runtime, however, the differences were not significant. The relevant aspect in terms of hardware is memory: experiences made with MSP430 and ARM Cortex M3 based hardware platforms revealed that the limiting factor is RAM capacity. Our experiences made during the trials show that problems typical for low-power and lossy networks can be addressed by the chosen network stack of CoAP, UDP, 6LoWPAN and 802.15.4; while still being lightweight enough to drive the application on the constrained devices investigated.
Citation	Bauer, R., Staudemeyer, C., Pöhls, H. C., and Fragkiadakis, A. 2016. “ECDSA on Things: IoT Integrity Protection in Practise,” *in Proceedings of the International Conference on Information and Communication Systems (ICICS 2016)*, Irbid, Jordan.
URL	https://link.springer.com/chapter/10.1007/978-3-319-50011-9_1

Marktscheffel et al. (2016): QR Code Based Mutual Authentication Protocol for Internet of Things	
Abstract	In the Internet of Things (IoT), security is important and challenging; however, it is often neglected. This paper presents a smart home scenario, together with its requirements for a secure and user friendly mutual authentication protocol. Protocols developed for the internet are often not applicable to the Internet of Things due to hardware limitations and physical inaccessibility of devices. To tackle the challenge of a usable and secure device authentication in the area of the IoT, a QR code based mutual authentication protocol is proposed. The protocol supports two operation modes to handle different hardware configurations with respect to cameras and displays. Both operation modes are secure against attacks within the proposed attacker model. The protocol can also be used to exchange the public keys between two parties, in order to establish a secure channel without a trusted third party.
Citation	Marktscheffel, T., Gottschlich, W., Popp, W., Werli, P., Fink, S. D., Bilzhause, A., and de Meer, H. 2016. “QR Code Based Mutual Authentication Protocol for Internet of Things,” in *Proceedings of the 5th Workshop on IoT-SoS: Internet of Things Smart Objects and Services (WOWMOM SOS-IOT 2016).*
URL	http://ieeexplore.ieee.org/document/7523562/

3.3 TP3 – Security Concepts for Virtualized Infrastructures

3.3.1 Project Overview

This subproject aims at investigating novel security concepts for virtualization-based environments, including cloud infrastructures and multiple interconnected clouds. The work plan of TP3 has been structured into two work packages: WP2, which addresses the detection and analysis of incidents, and WP1, which focuses on trustworthy execution. With the goal of enhancing the means for evidence collection and in-depth incident analysis, TP3 investigates methods and technologies at the hypervisor level, in particular making use of Virtual Machine Introspection (VMI). It analyzes the applicability of these methods in private cloud infrastructures, in which an investigator has full access to the cloud management infrastructure, as well as in public cloud infrastructures, in which such access is not feasible. TP3 also investigates how to correctly and efficiently handle virtual machine migration while a virtual machine is migrated in the cloud infrastructure, such that continuity of the monitoring can be assured in the transition from source to destination. Furthermore, TP3 addresses the problem of trustworthy execution in cloud environments as well as in federated cloud-of-cloud architectures. Combined with the incident detection and analysis approaches, monitoring data and analysis results shall be protected against malicious manipulation, including insider attacks at a cloud provider.

3.3.2 Results Achieved

During project execution, it became evident that the work on detection and analysis (WP2) needs to be considered the primary pillar of TP3, as it enables close collaboration with several of the other FORSEC subprojects. The main work on trustworthy execution (WP1) was delayed until a later project phase, as it ideally complements the detection and analysis as a means for supporting the integrity and probative value of collected incident data.

Based on a preliminary discussion of requirements, the project has defined several building blocks for enhancing incident investigations in cloud environments. The first achievement of TP3 is the definition of a modular low-level architecture for incident analysis based on virtual machine introspection (VMI), which enables the use of customizable analysis modules (Vlad and Reiser 2014).

The use of VMI on virtualization platforms on dedicated hosts is a well-established approach, but we faced the challenge of enabling the use of VMI in cloud computing environments. Several technical contributions of TP3 are important steps for enabling VMI in private and public cloud infrastructures.

Initially, we explored the feasibility of augmenting the cloud management layer with support for VMI-based host introspection. The LiveCloudInspector architecture and its prototype implementation for the OpenNebula cloud management platform integrate network forensics and remotely controlled VMI operations into a cloud management infrastructure (Zach and Reiser 2015; Zach 2014). We were able to show that our approach outperforms approaches proposed by other researchers and offers better flexibility.

State-of-the-art VMI approaches assume that the VMI operations are executed in a privileged system entity such as the Dom0 of the bare-metal hypervisor Xen or the host OS of a hosted hypervisor. During our work, we identified two challenges associated with such an approach. First, in a multi-tenant cloud, a strict separation of tenants is required. The LiveCloudInspector approach achieves this goal with custom access-control checks, but has the disadvantage of offering only a dedicated pre-defined set of VMI operations, instead of allowing arbitrary user-defined analysis. Second, the standard VMI approach is faced with the risk that the VMI-based analysis tool can be attacked and compromised by an adversary. The CloudPhylactor architecture (Taubmann, Rakotondravony, et al. 2016) is our contribution to tackling this problem. It is an important building block towards making our VMI-based analysis architecture integratable into public cloud infrastructures and making VMI more secure to use in both public and private clouds. CloudPhylactor proposes the use of mandatory access control to grant dedicated monitoring virtual machines the access to other selected virtual machines to perform VMI based operations.

In collaboration with other subprojects of the CLOUD cluster, our VMI approach was integrated into the CloudIDEA architecture (Fischer et al. 2015; Taubmann, Reiser, et al. 2015), which is based on a joint cluster scenario. The VMI-based monitoring with lightweight (suitable for continuous monitoring of a production system) and heavyweight (activated on demand for detailed incident analysis) tracing and analysis modules links to dynamic management of policies (TP4), anomaly detection based on machine learning (TP1), malware detection and anti-forensics (TP5) as well as economic evaluation of monitoring costs (TP10). A core concept of the cluster architecture is a decision engine that automatically controls what low-level detection and analysis methods to apply and how to react to suspicious activities, taking into account policies, cost models, and QoS constraints.

Another deficiency detected in discussions with other subprojects is the lack of detailed understanding of cloud-specific malware attacks. With dedicated focus on IaaS cloud infrastructures, we have systematically characterized and classified malware attacks that target virtual machines (Rakotondravony, Taubmann, et al. 2017) and reviewed recent publications in that area.

Our core architecture for VMI-based monitoring in cloud environments has been validated with evaluations of several use cases, including TLS monitoring, SSH honeypot implementation, visualization of malicious activities, and advanced malware detection.

- TLSkex (Taubmann, Frädrich, et al. 2016) (named TLSinspector (Taubmann, Dusold, et al. 2015) in an initial workshop publication) enables examining the content of encrypted TLS communication by extracting cryptographic session keys from VM memory using VMI and makes the content of such communication available for incident analysis and evidence collection. Compared to popular man-in-the-middleware proxy approaches for TLS monitoring, our approach does not tamper with the end-to-end connection, thus causing less impact on security, and is suitable for monitoring and analyzing applications that use TLS with certificate pinning.
- VMI-based monitoring can be used for efficiently implementing full-interaction honeypots that are stealthier (i.e., harder to distinguish from a real system) than established medium-interaction honeypots. We have implemented and discussed such an architecture for an SSH honeypot that collects detailed information about an attacker accessing a target system using VMI (Sentanoe 2017; Sentanoe et al. 2017a, 2017b).
- We presented a generic architecture that enables human users to visualize data acquired using VMI-based techniques for malware analysis purposes in IaaS clouds. The visualization architecture incorporates interactivity and control, allowing the user to trigger or adjust VMI monitoring operations on-demand for tailored and adaptive analysis (Rakotondravony and Reiser 2016; Rakotondravony, Köstler, et al. 2017).
- To evaluate the applicability of VMI for a cloud-based intrusion detection system we set up a test environment to gather execution traces of malware in a cloud environment. Together with TP1 we use these traces and train a classifier with the goal to build a tool that detects malware in a cloud system using system call based execution traces (Taubmann and Reiser 2016).

An inherent feature of many cloud environments is live migration of virtual machines, for example used for automated load balancing within a cloud data center. Migration entails both challenges and opportunities for VMI-based analysis. We have investigated the use of existing hypervisor-based migration mechanism for analyzing the VM memory during migration (Huppert 2015). Such approach is less powerful than what is feasible with state-of-the-art VMI techniques (for example, it enables static memory analysis, but not dynamic system call tracing of a running system). Nevertheless, it is a suitable option for providing some introspection mechanism in systems that do not offer a full VMI interface. In on-going work (Böhm 2017), we are extending the CloudPhylactor architecture to support more than one physical cloud node. We investigate the impact of live migration on concurrent VMI-based analysis. We aim at defining, implementing and validating an architecture that coordinates introspection and migration such that a migrating VM can continuously be monitored. This includes an evaluation of various mechanisms to perform VMI when migration is performed as well as the analysis of the corresponding access control mechanisms.

Besides supporting the detection and analysis of incidents, VMI-based approaches also enable the acquisition of forensic evidence in cloud environments. In collaboration with TP5, we have developed an abstract model (Freiling et al. 2017) that can be used for collecting forensic evidence in architectures composed of multiple abstraction layers (such as disk storage in the cloud and virtual memory in hypervisor-based nodes).

Our work on trustworthy execution (WP1) complements the results on the actual monitoring and analysis (WP2).

An initial analysis on TPM-based trusted computing in OpenNebula and confidentiality-preserving intrusion-tolerant architectures has yielded first insights into models for trust relations and a simple prototype for integrating TPM-based trusted computing into the OpenNebula cloud environment (Dawaras 2015).

We also collaborated with the associated FORSEC industry partner Fraunhofer AISEC and gained insight to the security architecture of trusted computing on mobile platforms and proposed a novel method of performing forensic analysis on them (Huber et al. 2016; Taubmann, Huber, et al. 2015).

The investigation of recent trusted execution technology (ARM Trustzone, Intel SGX) has led to external collaboration with Prof. Correia at INESC-ID, Portugal. The collaboration resulted in an on-going joint MSc thesis (Guerra 2017) as well as in a successful grant application for mutual visits (Bayerische Forschungsallianz, Intrusion Detection and Analysis in Clouds of Clouds, 2016), promoting the elaboration of an H2020 project proposal, planned to be submitted in August 2017 in the call DS-07-2017.

With the focus set on security monitoring in multiple interconnected cloud environments, we have analyzed the potential of applying resilience methods (intrusion-tolerant replication) developed in a concurrent local project

(OptScore, funded by DFG). We have defined an architecture for protecting monitoring information collected on nodes distributed on multiple independent cloud infrastructures (Reiser 2017). The implementation and validation of this approach is work in progress and will provide final results by end of the FORSEC project.

3.3.3 Contribution to FORSEC Research Alliance

TP3 contributes its research results in the area of trusted execution and incident analysis in virtualization environments in the area of interconnected cloud environments to the knowledge pool of CLOUD.

The CLOUD cluster collaborated on the definition of the integrated CloudIDEA (Cloud Intrusion Detection, Evidence preservation, and Analysis) architecture (Fischer et al. 2015; Taubmann, Reiser, et al. 2015). In close collaboration with TP10 and TP4, models, metrics and interfaces for using detection and analysis results in a policy-aware decision-making process have been defined. TP3 contributes the component design and implementations for virtual machine introspection to the integrated architecture.

Collaboration with TP1, TP4, TP5 and TP10 has been established for classifying malware attacks in IaaS cloud environment (Rakotondravony, Taubmann, et al. 2017). The objective is to provide a comprehensive analysis and classification of malware attacks that directly involve virtual machines. The work is completed by the analysis of the financial impact of malware attacks on real life businesses, as reported by existing literature.

We have developed a close collaboration for analyzing malware, based on our VMI monitoring and TUM's machine learning expertise. We were working on a joint architecture for running malware in a cloud-based VMI sandbox and analyzing the gathered traces (Taubmann and Kolosnjaj 2017).

Data collected with virtual machine introspection in cloud environments is valuable not only for malware detection and analysis, but also for forensic investigation. Forensic data acquisition in cloud environments is a field that has only recently moved into the focus of scientific research. In collaboration with TP5, we have develop an abstract model for collecting forensic traces in architectures composed of multiple layers as typically present for disk storage as well as for main memory in cloud infrastructures (Freiling et al. 2017).

Trustworthiness of data is relevant not only for forensic traces collected in cloud environments, but also for data originating in decentralized sensor nodes and processed by service platforms in Internet-of-Things environments, which is a core focus of TP2. We have collaborated with TP2 and external partners in the definition of an architecture that protects such sensor data in IoT devices and cloud nodes with redactable signatures, resulting in a joint publication (Frädrich et al. 2016).

3.3.4 Beyond FORSEC

As a practical output of FORSEC TP3, the CloudPhylactor implementation has proven exceptionally useful for many other activities. It is in active use in lab classes and student projects, as well as in most of our on-going research activities that make use of VMI.

Additionally, the gained insights and the *libvmtrace* implementation – a library for network and system call tracing based on VMI created for Tlskex and CloudPhylactor – are reused for the BMBF-funded project Dingfest in order to build a VMI-based intrusion detection system and to extract forensic evidence of virtual machines in cooperation with industry partners. We also plan to extend this library for the use on mobile platforms.

The development of a resilient architecture for collecting security monitoring data in interconnected cloud environments (Reiser 2017) not only contributes to the trustworthiness of that monitoring data. The architecture also provides a sample use case for our on-going OptScore project (funded by DFG) and will serve as use-case in a DFG grant application for a successor project (OptScore 2).

We also plan future research in the area of VMI on mobile devices, i.e., ARM based systems. Therefore, we want to explore how the trusted computing platform Trustzone of ARM can be used for trustworthy VMI. We successfully obtained a cooperation grant from Bavarian Research Alliance (BayIntAn grant IDACC) and plan to submit an EU project proposal (H2020 call DS-07-2017). Additionally, we have a joint on-going MSc thesis in which a student from Institute Superior Técnico in Portugal supervised by both Passau and Lisbon (Guerra 2017).

In the future, we plan more research in the fields of virtual machine introspection. We have submitted a DFG proposal "ARADIA: Cross-platform architecture for user-centric static and dynamic virtual machine introspection" which is currently under review. It aims to extend low-level VMI tracing techniques and to visualize the output for human operators.

3.3.5 Publications

Frädrich et al. (2016): Integrity and Authenticity Protection with Selective Disclosure Control in the Cloud and IoT	
Abstract	RSSRSS allow the redaction of parts from signed data. Updatable RSSRSS additionally enable the signatory to add new elements, while signatures can be merged by third parties under certain conditions. We propose a framework for two new real-life application scenarios and implement it using an RSSRSS with sufficient functionality on three different platforms, ranging from a potent cloud to a very resource-constrained Android device. Our evaluation shows impractical run time especially on the IoT device for the existing construction that was proven to be secure in the standard model. Thus, we provide an adjusted scheme with far better performance, which we prove to be secure in the random oracle model. Furthermore, we show how to increase performance using parallelization and several optimizations.
Citation	Frädrich, C., Pöhls, H. C., Popp, W., Rakotondravony, N., and Samelin, K. 2016. "Integrity and Authenticity Protection with Selective Disclosure Control in the Cloud and IoT," in *Proceedings of the International Conference on Information and Communication Systems (ICICS 2016),* Irbid, Jordan.
URL	https://link.springer.com/chapter/10.1007/978-3-319-50011-9_16

Huber et al. (2016): A Flexible Framework for Mobile Device Forensics Based on Cold Boot Attacks	
Abstract	Mobile devices, like tablets and smartphones, are common place in everyday life. Thus, the degree of security these devices can provide against digital forensics is of particular interest. A common method to access arbitrary data in main memory is the cold boot attack. The cold boot attack exploits the remanence effect that causes data in DRAM modules not to lose the content immediately in case of a power cut-off. This makes it possible to restart a device and extract the data in main memory. In this paper, we present a novel framework for cold boot-based data acquisition with a minimal bare metal application on a mobile device. In contrast to other cold boot approaches, our forensics tool overwrites only a minimal amount of data in main memory. This tool requires no more than three kilobytes of constant data in the kernel code section. We hence sustain all of the data relevant for the analysis of the previously running system. This makes it possible to analyze the memory with data acquisition tools. For this purpose, we extend the memory forensics tool Volatility in order to request parts of the main memory dynamically from our bare metal application. We show the feasibility of our approach on the Samsung Galaxy S4 and Nexus 5 mobile devices along with an extensive evaluation. First, we compare our framework to a traditional memory dump-based analysis. In the next step, we show the potential of our framework by acquiring sensitive user data.
Citation	Huber, M., Taubmann, B., Wessel, S., Reiser, H. P., and Sigl, G. 2016. "A Flexible Framework for Mobile Device Forensics Based on Cold Boot Attacks," EURASIP Journal on Information Security.
URL	https://link.springer.com/article/10.1186/s13635-016-0041-4

Rakotondravony, Köstler, et al. (2017): Towards a Generic Architecture for Interactive Cost-aware Visualization of Monitoring Data in Distributed Systems	
Abstract	The collection of monitoring data in distributed systems can serve many different purposes, such as system status monitoring, performance evaluation, and optimization. There are many well-established approaches for data collection and visualization in these areas. For objectives such as debugging complex distributed applications, in-depth analysis of malicious attacks, and forensic investigations, the joint analysis and visualization of a large variety of data gathered at different layers of the system is of great value. The utilization of heavy-weight monitoring techniques requires a cost-aware on demand activation of such monitoring. We present an architecture for an interactive and cost-aware visualization of monitoring data combined from multiple sources in distributed systems. We introduce two distinguishing properties: the possibilities to reconfigure data collection and a cost prediction mechanism that supports the user in a cost-aware, dynamic activation of monitoring components in an interactive in-depth analysis. We illustrate the use of such cost prediction for monitoring using VMI-based mechanisms.
Citation	Rakotondravony, N., Köstler, J., and Reiser, H. P. 2017. "Towards a Generic Architecture for Interactive Cost-aware Visualization of Monitoring Data in Distributed Systems," in *Proceedings of the 17th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS 2017),* Neuchatel, Switzerland.

Rakotondravony and Reiser (2016): Visualizing and Controlling VMI-Based Malware Analysis in IaaS Cloud	
Abstract	Security in virtualized environment has known the support of different tools in the low-level detection and analysis of malware. The in-guest tracing mechanisms are now capable of operating at assembly language-, system call-, function call-and instruction-level to detect and classify malicious activities. Therefore, they are producing large amount of data about the state of a target system. However, the integrity of such data becomes questionable whenever the hosting target system is compromised. With virtual machine introspection (VMI), the monitoring tool runs outside the target monitored virtual machine (VM). Thus, the integrity of retrieved data is ensured even if the target system is compromised. Various works have brought VMI to Infrastructure-as-a-Service (Iaas) cloud environment, allowing the cloud user to run (simultaneous) forensics operations on his production VMs. The associated tracing mechanisms can collect larger amount of data in form of commented behavior traces or unstandardized log records. Thus, a human operator is needed to efficiently parse, represent, visualize and interpret the collected data, to benefit from their security relevance. The use of visualization helps analysts investigate, compare and culster malware samples. Existing visualization tools make use of recorded information to enhance the detection of intrusive behavior or the clustering of malware from the observed system. However, at our knowledge, no existing tools establish a pre-to post-exploitation visualization graphs. We present an approach that enhances the detection and analysis of malware in the cloud by providing the cloud end-users the mean to efficiently visualize the different security relevant data collected through multiple VMI-based mechanisms.
Citation	Rakotondravony, N., and Reiser, H. P. 2016. "Visualizing and Controlling VMI-based Malware Analysis in IaaS Cloud," in Symposium on Reliable Distributed Systems (SRDS 2016), PhD Forum.
URL	http://ieeexplore.ieee.org/document/7794347/

Reiser (2017): Towards Intrusion-resilient Security Monitoring in Multi-cloud Infrastructures	
Abstract	Multi-cloud architectures enable the design of resilient distributed service applications. Such applications can benefit from a combination of intrusion-tolerant replication across clouds with intrusion detection and analysis mechanisms. Such mechanisms enable the detection of attacks that affect multiple replicas and thus exceed the intrusion masking capability, and in addition support fast reaction and recovery from local intrusions. In this work-in-progress paper we present a security analysis on which an intrusion detection and analysis service can be based on. We sketch the architecture of such a cross-cloud intrusion detection architecture that combines a set of well-known mechanisms. The goal of our approach is obtaining a resource-efficient service with optimal resilience against malicious attacks.
Citation	Reiser, H. P. 2017. "Towards Intrusion-resilient Security Monitoring in Multi-cloud Infrastructures," in Workshop on Security and Dependability of Multi-Domain Infrastructures, EuroSys 2017, Belgrade, Serbia.
URL	https://dl.acm.org/citation.cfm?id=3071067

Sentanoe et al. (2017a): Virtual Machine Introspection Based SSH Honeypot	
Abstract	A honeypot provides information about the new attack and exploitation methods and allows analyzing the adversary's activities during or after exploitation. One way of an adversary to communicate with a server is via secure shell (SSH). SSH provides secure login, file transfer, X11 forwarding, and TCP/IP connections over untrusted networks. SSH is a preferred target for attacks, as it is frequently used with password-based authentication, and weak passwords are easily exploited using brute-force attacks. In this paper, we introduce a Virtual Machine Introspection based SSH honeypot. We discuss the design of the system and how to extract valuable information such as the credential used by the attacker and the entered commands. Our experiments show that the system is able to detect the adversary's activities during and after exploitation, and it has advantages compared to currently usedSSH honeypot approaches.
Citation	Sentanoe, S., Taubmann, B., and Reiser, H. P. 2017b. “Virtual Machine Introspection Based SSH Honeypot,” in *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems (SHCIS 2017),* Neuchâtel, Switzerland.

Taubmann, Dusold, et al. (2015): Analysing Malware Attacks in the Cloud: A Use Case for the TLSInspector Toolkit	
Abstract	Nowadays, malicious attacks in the Internet often use encrypted communication channels. Thus, an attacker might exploit a vulnerability in a web service using the HTTPS protocol. If network intrusion detection systems (NIDS) are unable to decrypt this communication, they cannot observe the contents of such attacks. If the NIDS is operated independently of the web services, it is impractical to directly provide decryption keys to it. This is, for example, the case if a cloud provider operates the NIDS, while a cloud customer manages the web service within a virtual machine. Additionally, malware often encrypts the communication to a command and control server. The encryption keys used for that communication channel are fully under the control of the malware and thus it is even more difficult to provide them to the NIDS. This paper discusses both use cases in a common cloud scenario and describes a VMI based prototype that is able to decrypt TLS encrypted communication of a virtual machine. The decryption is achieved by taking a memory snapshot and extracting the cryptographic key that is required to decrypt a network flow. We experimentally evaluate the overhead caused by taking the memory snapshots and the performance of extracting the encryption key from the snapshot.
Citation	Taubmann, B., Dusold, D., Frädrich, C., and Reiser, H. P. 2015. “Analysing Malware Attacks in the Cloud: A Use Case for the TLSInspector Toolkit,” in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015),* Vienna, Austria.

Taubmann, Huber, et al. (2015): A Lightweight Framework for Cold Boot Based Forensics on Mobile Devices	
Abstract	Mobile devices, like tablets and smartphones, are common place in everyday life. Thus, the degree of security these devices can provide against digital forensics is of particular interest. A common method to access arbitrary data in main memory is the cold boot attack. The cold boot attack exploits theremanence effect that causes data in DRAM modules not to lose the content immediately in case of a power cut-off. This makes it possible to restart a device and extract the data in main memory. In this paper, we present a novel framework for cold boot based data acquisition with a minimal bare metal application on a mobile device. In contrast to other cold boot approaches, our forensics tool overwrites only a minimal amount of data in main memory. This tool requires no more than five kilobytes of constant data in the kernel code section. We hence sustain all of the data relevant for the analysis of the previously running system. This makes it possible to analyze the memory with data acquisition tools. For this purpose, we extend the memory forensics tool Volatility in order to request parts of the main memory dynamically from our bare metal application. We show the feasibility of our approach by comparing it to a traditional memory dump based analysis using the Samsung Galaxy S4 mobile device.
Citation	Taubmann, B., Huber, M., Heim, L., Sigl, G., and Reiser, H. P. 2015. “A Lightweight Framework for Cold Boot Based Forensics on Mobile Devices,” in *Proceedings of the 10th International Conference on Availability, Reliability and Security (ARES 2015)*, Toulouse, France.
URL	http://ieeexplore.ieee.org/document/7299905/

Taubmann, Frädrich, et al. (2016): TLSkex: Harnessing Virtual Machine Introspection for Decrypting TLS Communication	
Abstract	Nowadays, many applications by default use encryption of network traffic to achieve a higher level of privacy and confidentiality. One of the most frequently applied cryptographic protocols is Transport Layer Security (TLS). However, also adversaries make use of TLS encryption in order to hide attacks or command & control communication. For detecting and analyzing such threats, making the contents of encrypted communication available to security tools becomes essential. The ideal solution for this problem should offer efficient and stealthy decryption without having a negative impact on over-all security. This paper presents TLSkex (TLS Key EXtractor), an approach to extract the master key of a TLS connection at runtime from the virtual machine's main memory using virtual machine introspection techniques. Afterwards, the master key is used to decrypt the TLS session. In contrast to other solutions, TLSkex neither manipulates the network connection nor the communicating application. Thus, our approach is applicable for malware analysis and intrusion detection in scenarios where applications cannot be modified. Moreover, TLSkex is also able to decrypt TLS sessions that use perfect forward secrecy key exchange algorithms. In this paper, we define a generic approach for TLS key extraction based on virtual machine introspection, present our TLSkex prototype implementation of this approach, and evaluate the prototype.
Citation	Taubmann, B., Frädrich, C., Dusold, D., and Reiser, H. P. 2016. "TLSkex: Harnessing virtual machine introspection for decrypting TLS communication," in DFRWS EU 2016 Annual Conference.
URL	http://www.sciencedirect.com/science/article/pii/S1742287616300081

Taubmann et al. (2016): Harnessing Mandatory Access Control for Virtual Machine Introspection in Cloud Data Centers	
Abstract	Virtual machine introspection is a valuable approach for malware analysis and forensic evidence collection on virtual machines. However, there are no feasible solutions how it can be used in production systems of cloud providers. In this paper, we present the CloudPhylactor architecture. It harnesses the mandatory access control of Xen to grant dedicated monitoring virtual machines the rights to access the main memory of other virtual machines in order to run introspection operations. This allows customers to create monitoring virtual machines that have access to perform VMI-based operations on their production virtual machines. With our prototype implementation, we show that our approach does not introduce performance drawbacks and gives cloud customers full control to do introspection on their virtual machines. We also show that the impact of successful attacks to the monitoring framework is reduced.
Citation	Taubmann, B., Rakotondravony, N., and Reiser, H. P. 2016. “CloudPhylactor: Harnessing Mandatory Access Control for Virtual Machine Introspection in Cloud Data Centers,” in The 15th IEEE International Conference on Trust, Security and Privacy in Computing and Communications (IEEE TrustCom-16).
URL	http://ieeexplore.ieee.org/document/7847045/

Taubmann and Kolosnjaj (2017): Architecture for Resource-Aware VMI-based Cloud Malware Analysis	
Abstract	Virtual machine introspection (VMI) is a technology with many possible applications, such as malware analysis and intrusion detection. However, this technique is resource intensive, as inspecting program behavior includes recording of a high number of events caused by the analyzed binary and related processes. In this paper we present an architecture that leverages cloud resources for virtual machine-based malware analysis in order to train a classifier for detecting cloud-specific malware. This architecture is designed while having in mind the resource consumption when applying the VMI-based technology in production systems, in particular the overhead of tracing a large set of system calls. In order to minimize the data acquisition overhead, we use a data-driven approach from the area of resource-aware machine learning. This approach enables us to optimize the trade-off between malware detection performance and the overhead of our VMI-based tracing system.
Citation	Taubmann, B., and Kolosnjaj, B. 2017. “Architecture for Resource-Aware VMI-based Cloud Malware Analysis,” in *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems (SHCIS 2017),* Neuchâtel, Switzerland.

Vlad and Reiser (2014): Towards a Flexible Virtualization-based Architecture for Malware Detection and Analysis	
Abstract	The complexity and sophistication of malicious attacks against IT systems have steadily increased over the past decades. Tools used to detect and analyse such attacks need to evolve continuously as well in order to cope with such attacks. In this paper, we identify some limitation of existing approaches and propose a novel architecture for an attack detection and analysis framework. This architecture is based on virtualization technology to execute target systems, supports a broad spectrum of low-level tracing modules and sophisticated, extensible virtual-machine introspection mechanisms, combined with an extensible plug-in interface for specialized detection and analysis mechanisms, and it offers support for deployment in cloud infrastructures.
Citation	Vlad, M., and Reiser, H. P. 2014. "Towards a Flexible Virtualization based Architecture for Malware Detection and Analysis," in *Proceedings of the 1st Workshop on Security in Highly Connected IT Systems (SHCIS 2014*), Munich, Germany.
URL	http://ieeexplore.ieee.org/document/6974866/

Zach and Reiser (2015): LiveCloudInspector. Towards Integrated IaaS Forensics in the Cloud	
Abstract	Cloud-based systems are becoming an increasingly attractive target for malicious attacks. In IaaS environments, malicious attacks on a cloud customer's virtual machine may affect the customer, who cannot use all diagnostic means that are available in dedicated in-house infrastructures, as well as the cloud provider, due to possible subsequent attacks against the cloud infrastructure and other co-hosted customers. This paper presents an integrated approach towards forensics and incident analysis in IaaS cloud environments. The proposed architecture enables the cloud provider to securely offer forensics services to its customers on a self-service platform. The architecture combines three important analysis techniques and provides significantly better investigation capabilities than existing systems: First, it supports host-based forensics based on virtual machine introspection. Second, it offers live remote capture of network traffic. Third, and most importantly, it provides hybrid combinations of the first two techniques, which enables enhanced analysis capabilities such as support for monitoring encrypted communication.
Citation	Zach, J., and Reiser, H. P. 2015. "LiveCloudInspector: Towards Integrated IaaS Forensics in the Cloud," in *Proceedings of the 15th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS 2015),* Grenoble, France.
URL	https://link.springer.com/chapter/10.1007/978-3-319-19129-4_17

3.4 TP4 – Secure Migration of Virtual Machines

3.4.1 Project Overview

The main research problem in TP4 is autonomic security management in IaaS clouds, in which the customer deploys a Virtual Network (VN) consisting of multiple Virtual Machine (VMs) to realize a certain service. The main goals are early reaction to threats through VN reconfiguration, and secure VM live migration required to realize the reconfiguration. These goals are approached by supporting the security measurements with a decision making architecture that considers on the one hand the threats and events in the environment, and on the other hand the SLA between the cloud provider and customer. A Decision Engine (DE) has been designed and the significant SLA metrics and relevant VM placement policies have been identified and extended to cover VNs deployed in the cloud. The security-aware placement polices have been realized using Virtual Network Embedding (VNE) algorithms. Moreover, the project has designed an architecture for SLA-aware secure live migration of VMs, and identified the threats and costs imposed by the migration. Furthermore, the project has contributed to the classification of the IaaS cloud-specific attacks, and defining possible early detectable behavior patterns.

3.4.2 Results Achieved

Secure Migration

Live migration transfers a running VM to another host with as little interruption (downtime) as possible. It supports the continuity of services under maintenance, faults, and attack conditions. Moreover, it can improve the operating costs and performance via VMs consolidation and load-balancing, respectively. Live migration opens new threats in IaaS clouds, in particular when migrating VMs between different data centers. Two main threats imposed by live migration are the exploitation of the migration itself, and attacks on the customer's VM during migration. An example of the first threat is forcing the cloud management system to create many migrations, leading to a DOS attack on the VMs and hosts. This can be performed for example by varying the resource utilization of malicious VMs to trigger live migration. An example of the second threat is the man-in-the-middle attack that either changes the migrated data or eavesdrops on the VM to extract sensitive data such as passwords and keys. Since securing live migration is a serious issue, secure migration mechanisms are required. However, these mechanisms cause additional overhead.

The threats, security requirements, and security mechanisms of live migration have been surveyed. There is little or no consideration of the QoS cost, and there is no comprehensive SLA-aware solution. Our research problem is to flexibly secure the migration and investigate its costs. In our proposed solution concept, the secure migration algorithm considers the security threats of the migration path, and the SLA metrics of the VN to force certain security mechanisms. To this end, the main service metric under focus is the downtime. To use this SLA availability metric, the solution proposes to use a downtime budget for every VM in the VN. Live migration downtime is currently being modeled with the main parameters: migration strategy, network bandwidth, VM's memory size/utilization and CPU load, and the encryption mechanism used. A first model related to VM memory utilization has been built and is currently being validated.

A testbed has been built using OpenNebula virtual environment (with KVM hypervisor), in which SSH is used for live migration. An architecture for measuring the migration downtime in OpenNebula has been built. The VMs have been outfitted with networking tools allowing them to communicate to external machines via SSH and other methods. A simple tool has been implemented in order to measure the migration downtime. The tool consists of two parts, the server runs on the VM that is to be migrated, and the client runs on any external machine. When launched, the client communicates with OpenNebula in order to initiate the migration of the VM and sends the server running on the VM a signal. Once the server receives the signal, it sends the client small TCP packets every ten milliseconds. The client keeps track of the delay between the arrival times of these packets and periodically checks in with OpenNebula to verify the status of the migration. When the VM goes down during the last phase of live migration, there is a noticeable break in the stream of packets the client receives from the server. The aforementioned model parameters are carefully controlled during the experiments. Each test performs 100 concurrent live migrations and measures their total migration time and downtime.

Decision Engine

A DE for security management in IaaS clouds has been designed, and a prototype has been developed. The DE receives events from an assumed cloud monitoring system that deploys lightweight tracing mechanisms to monitor VM activities in the production environment. The events represent certain suspected attacks based on pre-defined suspicious behavior patterns. An example pattern is the memory usage spike for more than 5 seconds by the *Kelihos* malware that performs a DDOS attack.

Behavior pattern: *{Pattern name, Parameter, Suspected malware, Suspected attack}*
Example: *{Memory_usage_spike, period > 5 seconds, Kelihos, DDOS}*

The DE reacts to events according to user-defined policies. The event defines the suspected VM and suspected attack/malware. We assume that the provider needs individual policies for each VM according to its SLA. The policy defines a set of reactions to a pre-defined event. The main threat reaction under focus in the project is isolating a suspicious VM in a dedicated analysis environment (protected host) to protect the production environment from attacks and performance degradation. The DE uses a migration algorithms to select the protected host that do not host any VM and offers enough computing resources to host the target VM. The following shows the proposed decision structures for the DE with examples.

Event: *{VMID, Malware, Attack}* - Example: *{VM 100, Kelihos, DDOS}*
Policy: *{Malware, Attack, VMID, Actions}* - Example: *{-, DDOS, VM 100, migrate-stop}.*

The DE parses the user-defined policies from a policy fie using XML format. The policy defines multiple action sets per event. All action sets related to a certain event are executed (in parallel). The DE tries the actions in a set sequentially and only the first possible action is executed. In the following, an example policy is presented. If VM 100 is showing a behavior pattern that refers to a DDOS attack, the DE first tries to migrate it to a protected host. If this action fails, the VM is stopped.

<Policy> <Event> <VM>100</VM><Malware> - </ Malware > < Attack > DDOS </ Attack></Event>

<ActionSets><Set> <Action>migrate</Action><Action>stop</Action> </Set></ActionSets> </Policy>

The DE prototype is provided with a generic interface to cloud APIs. A communication interface to each API type (such as RPC/XML in OpenNebula) is required to parse the functions needed to monitor and configure the environment. The interface uses an environment-specific driver that defines the required API functions with their parameters and returned data structures. The interface and driver for OpenNebula have been developed. The following example sketches some important functions from the OpenNebula driver: the API information, VM status, and VM stop functions. The DE parses the functions (using the RPC/XML interface) and replaces the missing parameters (such as $VMID$) with the required values.

```
[API_Info]
API = XML-RPC
Server = http://$CloudIP$:2633/RPC2
 [VM]
VM_Status_URL = one.vm.info
VM_Status_Parameter = $username$:$password$,$VMID$|int
VM_Stop_URL = one.vm.action
```

VM_Stop_Parameter = $username$:$password$,poweroff-hard,$VMID$|int

All functional tests of the porotype have been successfully performed. The future work will integrate SLAs, network topology, and available security mechanisms in the DE.

Cloud SLAs and Placement Policies

QoS, privacy, and security are main challenges in migrating business operations to the cloud, which are covered only separately by researchers. The cloud provider needs placement policies for VNs and security mechanisms to conform to SLAs (to avoid the financial and reputational losses), protect its environment, and maximized its revenue. The customer might deploy a full service VN in the IaaS cloud, which might request different guaranties such as availability and data privacy. A comprehensive view of the aforementioned challenges through SLAs and placement policies is still missing. In the project, we have surveyed them and proposed new VN-level SLA metrics and the related VM/service placement policies. In the following, we list the main SLA metrics and placement policies discussed:

SLA metrics:

- QoS: Service availability and response time.
- Security: Co-location with other customers, offering security mechanisms such as firewalls.
- Privacy: Placement in certain geographical locations, permission for VM analysis.
- Financial: Service prices, penalties under violation of SLA metrics.
- Exceptional conditions for violating the SLA. For example, a suspicious behavior.

Deployment policies:

- QoS: Service dedicated hosts, backup VMs, placement in nearby locations.
- Security: Customer dedicated hosts, hosts/networks with certain security mechanisms.
- Privacy: Placement in certain locations.
- Provider policies: Analyzing/stopping a suspicious VM, prioritizing VMs according to the revenue/loss under threat or conflict of policies/SLAs.

Furthermore, we have identified some serious challenges of the problem of SLA management in IaaS clouds. For example, collecting evidences about suspicious behavior, how the customer checks the provider's compliance with the security/privacy SLA, and how the conflict between customer's privacy and the security of other customers and the cloud provider can be resolved. Finally, a use case that proves the need for such a comprehensive view of SLAs and service placement has been presented.

Security-aware VNE

The current approaches to security-aware VNE only address abstract security levels and metrics based on available security mechanisms, and do not consider concrete security mechanisms such as firewalls. In the project, we have classified, defined, and modeled a basic set of security requirements of VNs. The project has identified and modeled the topological constraints as a new type of constraints that requires additional support by VNE. A topology constraint affects an entire subnet. For example, the VN might specify network domains that should be separated. The cross-domain links in this case should be mapped through firewalls by the link mapping stage of VNE. Another topological constraint is that virtual domains must not be split by a firewall and must be mapped onto a single physical domain. In this model, mapping the security requirements of the virtual link should check for certain properties along its physical path. The typical VNE demand/resource model has been extended to model security capabilities and demands. We have modeled the typical security requirements (such as Trusted Hardware, Network Intrusion Detection System, and Firewall) as VNE node, link, and topological demands. Furthermore, we have provided a proof-of-concept implementation of this new security-aware VNE model in our VNE tool, ALEVIN, and incorporated the constraints into VNE algorithms.

3.4.3 Contribution to FORSEC Research Alliance

The main contribution of TP4 to FORSEC is designing a DE as a central coordinator in the CloudIDEA malware defense architecture proposed by Cluster 3 (CLOUD). The DE reacts to a possible indicator of an attack by initiating certain actions according to the expected attack, SLAs, and policies defined by the cloud provider. The most significant action is migrating a suspicious VM to a dedicated analysis environment, to protect the production environment from the suspicious VM and avoid overloading it with heavy analysis mechanisms. In this case, a full VN reconfiguration might be required to fulfill the customer's SLA. The interfaces, inputs, and outputs of the DE have been defined.

The second important contribution is classifying IaaS cloud malware. A literature survey about attack vectors, attack behavior, defense measurements, and technical reports about attacks has been conducted, and a first behavior pattern has been identified. Typical IaaS cloud-specific attacks are cross-VM cache-based side-channel attacks. In these attacks, the attacker needs first to gain and verify co-location with the victim VM. The methods of forcing and verifying co-location with a victim VM use VM instance types, availability zones, IP addresses, and ICMP. Those methods expect the behavior of VM placement algorithms and use also side-channels. For example, by applying a varying load on the victim VM and measuring the memory access time of the attacker's VM. The noise of other VMs and the host are challenges to these attacks. An example of defense measurements is VM location obfuscation using user-defined or dynamic VM placement algorithms.

A typical cache-based side-channel attack is Prime and Probe attack. The attack fills the cache then measures the memory access time. If the access time of a certain memory address is higher than a threshold, a memory page associated with the respective cache line was accessed by the victim. A researcher demonstrated a Prime and Probe attack used to recover a full 2048 bit RSA key in Amazon EC2 cloud. Another cache-based side-channel attack is Flush and Reload attack. The attack relies on memory deduplication feature (sharing common system libraries in the memory). It flushes certain known library memory lines from the cache, waits, then measures the access time. We have found a source code used to extract the RSA private key from a certain RSA library. We have reproduced the experiment in OpenNebula, but we couldn't extract the key. Frequent flushing of the cache is a candidate behavior pattern of such an attack.

The third contribution to FORSEC is a collaboration with TP10 that has surveyed the research about QoS, privacy, and security issues in IaaS clouds and the financial impacts of non-compliance to SLAs. The collaboration has defined comprehensive SLA metrics and placement policies that address these issues at the level of VNs deployed in an IaaS cloud.

3.4.4 Beyond FORSEC

Two significant challenges in IaaS clouds are how to achieve a comprehensive and autonomic SLA management, and how to flexibly deploy and activate security mechanisms when needed. A future research direction is integrating the DE, security-aware VNE model, and secure migration architecture, and provide them with the defined VN SLA metrics and placement policies. To address the flexible deployment of security mechanisms, Software-Defined Networking (SDN) and Network Function Virtualization (NFV) will be considered. These are modern networking trends, on which the future mobile networks (5G) and industrial networks (Industry 4.0) will be based. SDN offers a centralized network control, while NFV offers flexible deployment of network functions as VMs in a cloud. VNE coordinated SDN/NFV deployment is needed to satisfy the requirements of future networks such as high throughput, security, fault tolerance, and energy efficiency. In the NFV mapping problem, the combination of network functions imposes a non-fixed topology. Several variations of the topology might be possible, for example, encryption function before or after compression function? This has effects on resource usage and network security level. An example scenario of the NFV mapping problem is determining and adding virtualized firewalls to the VN before embedding it.

Another related future goal is using SDN/NFV for secure migration. The research questions are:

- Which network functions can be virtualized to support secure migration?
- Where to deploy these functions, and how to build communication services among them using SDN?

A possible scenario is virtualizing advanced encryption functions and deploying them as migration gateways in the source and destination data centers. The challenge in such a scenario is intercepting live migration, which is an environment-specific problem.

3.4.5 Publications

Alshawish et al. (2017): Playing a Multi-objective Spot-checking Game in Public Transportation Systems	
Abstract	Public transportation systems represent an essential sector of any nation's critical infrastructure. Hence, continuity of their services is deemed important and with a high priority to the nations. Concerns over risks like terrorism, criminal offenses, and business revenue loss impose the need for enhancing situation awareness in these systems. However, practices, such as conducting random patrols or regular spot-checks on passengers to prevent or deter potential violations, are strictly limited by the number of available resources (e.g. security staff or fare inspectors) and by the ability of potential opponents (e.g. criminals, or fare evaders) to predict or observe the inspectors' presence patterns. Casting the interactions between these competitive entities (inspectors/security offcials and criminals/fare dodgers) into a game-theoretic model will enable involved system operators to 1) find optimal cost-effective (or multi-goal) human resource allocation or spot-check schedules, 2) capture and treat uncertainty due to imperfectness of information, 3) integrate measurements from heterogeneous natures (e.g. statistics, expert opinions, or simulation results). This work applies a game-theoretical model that uses random probability-distribution-valued payoffs to allow playing spot-checking games with diverging actions' outcomes as well as avoiding information loss due to combining several measurements into one representative (e.g. average).
Citation	Alshawish, Alille; Abid, Mohamed Amine; Rass, Stefan; de Meer, Hermann (2017): Playing a Multi-objective Spot-checking Game in Public Transportation Systems. In: *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems (SHCIS 2017).* Neuchâtel, Switzerland.

Fischer et al. (2015): A Malware Defense Architecture for Cloud Data Centers	
Abstract	Due to the proliferation of cloud computing, cloud-based systems are becoming an increasingly attractive target for malware. In an Infrastructure-as-a-Service (IaaS) cloud, malware located in a customer's virtual machine (VM) affects not only this customer, but may also attack the cloud infrastructure and other co-hosted customers directly. This paper presents CloudIDEA, an architecture that provides a security service for malware defens in cloud environments. It combines lightweight intrusion monitoring with on-demand isolation, evidence collection, and in-depth analysis of VMs on dedicated analysis hosts. A dynamic decision engine makes on-demand decisions on how to handle suspicious events considering cost-efficiency and quality-of-service constraints.
Citation	Fischer, A., Kittel, T., Kolosnjaji, B., Lengyel, T. K., Mandarawi, W., Reiser, H. P., Taubmann, B., Weishäupl, E., de Meer, H., Mu, T., and Protsenko, M. 2015. “CloudIDEA: A Malware Defense Architecture for Cloud Data Centers,” in *Proceedings of the 5th International Symposium on Cloud Computing, Trusted Computing and Secure Virtual Infrastructures - Cloud and Trusted Computing (C&TC),* Rhodes, Greece.
URL	https://link.springer.com/chapter/10.1007/978-3-319-26148-5_40

Fischer and de Meer (2016): Generating Virtual Network Embedding Problems with Guaranteed Solutions	
Abstract	The efficiency of network virtualization depends on the appropriate assignment of resources. The underlying problem, called virtual network embedding, has been much discussed in the literature, and many algorithms have been proposed, attempting to optimize the resource assignment in various respects. Evaluation of those algorithms requires a large number of randomly generated embedding scenarios. This paper presents a novel scenario generation approach and demonstrates how to produce scenarios with a guaranteed exact solution, thereby, facilitating better evaluation of embedding algorithms.
Citation	Fischer, A., and de Meer, H. 2016. "Generating Virtual Network Embedding Problems with Guaranteed Solutions," IEEE Transactions on Network and Service Management (13:3), pp. 504–517.
URL	http://ieeexplore.ieee.org/document/7527632/

Fischer et al. (2016): Modeling Security Requirements for VNE algorithms	
Abstract	Public and private Infrastructure as a Service (IaaS) clouds are widely used by individuals and organizations to provision flexible virtual computing resources on demand. Virtual Network Embedding (VNE) algorithms are employed in this context to provide an automated resource assignment. With multiple involved parties, security-aware Virtual Machine (VM) placement becomes highly relevant for production environments. Moreover, VNE algorithms should also consider the security requirements of the interconnections between VMs, thereby extending the problem to networks. This paper discusses security requirements of Virtual Networks (VNs) and shows how they can be modeled in VNE to map them to the provided security mechanisms in the physical network. The paper also presents an implementation of this security-aware VNE model in the public simulation platform ALEVIN, demonstrating the applicability with a realistic use case of such a model.
Citation	Fischer, A., Kuehn, R., Mandarawi, W., and de Meer, H. 2016. “Modeling Security Requirements for VNE algorithms,” in *Proceedings of the 10th EAI International Conference on Performance Evaluation Methodologies and Tools.*

Mandarawi et al. (2015): QoS-Aware Secure Live Migration of Virtual Machines	
Abstract	The live migration of Virtual Machines (VMs) is a key technology in server virtualization solutions used to deploy Infrastructure-as-a-Service (IaaS) clouds. This process, on one hand, increases the elasticity, fault tolerance, and maintainability in the virtual environment. On the other hand, it increases the security challenges in cloud environments, especially when the migration is performed between different data centers. Secure live migration mechanisms are required to keep the security requirements of both cloud customers and providers satisfied. These mechanisms are known to increase the migration downtime of the VMs, which plays a significant role in the compliance to Service Level Agreements (SLAs). This paper discusses the main threats caused by live migration and the main approaches for securing the migration. The requirements of a comprehensive Quality of Service (QoS)-aware secure live migration solution that keeps both security and QoS requirements satisfied are defined.
Citation	Mandarawi, W., Fischer, A., de Meer, H., and Weishäupl, E. 2015. "QoS-Aware Secure Live Migration of Virtual Machines," in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015)*, Vienna, Austria.

Mandarawi et al. (2016): Constraint-Based Virtualization of Industrial Networks	
Abstract	In modern industrial solutions, Ethernet-based communication networks have been replacing bus technologies. Ethernet is no longer found only in inter-controller or manufacturing execution systems, but has penetrated into the real-time sensitive automation process (i.e., close to the machines and sensors). Ethernet itself adds many advantages to industrial environments where digitalization also means more data-driven IT services interacting with the machines. However, in order to cater to the needs of both new and more automation-related communication, a better restructuring of the network and resources among multitenant systems needs to be carried out. Various Industrial Ethernet (IE) standards already allow some localized separation of application flows with the help of Quality of Service (QoS) mechanisms. These technologies also expect some planning or engineering of the system which takes place by estimating worst-case scenarios of possible traffic generated by all assumed applications. This approach, however, lacks the flexibility to add new services or to extend the system participants on the fly without a major redesign and reconfiguration of the whole network. Network virtualization and segmentation is used to satisfy these requirements of more support for dynamic scenarios, while keeping and protecting time-critical production traffic. Network virtualization allows slicing of the real physical network connecting a set of applications and end devices into logically separated portions or Slices. A set of resource demands and constraints is defined on a Slice or Virtual Network level. Slice links are then mapped over physical paths starting from end devices through forwarding devices that can guarantee these demands and constraints. In this chapter, the modeling of virtual industrial network constraints is addressed with a focus on communication delay. For evaluation purposes, the modeled network and mapping criteria are implemented in the Virtual Network Embedding (VNE) traffic-engineering platform ALEVIN [1].
Citation	Mandarawi, W., Fischer, A., Houyou, A. M., Huth, H.-P., and de Meer, H. 2016. “Constraint-Based Virtualization of Industrial Networks,” in Principles of Performance and Reliability Modeling and Evaluation, pp. 567–586.
URL	https://link.springer.com/chapter/10.1007%2F978-3-319-30599-8_22

3.5 TP5 – Software Protection and Anti Forensics

3.5.1 Project Overview

In highly connected systems, executable programs are typically shared in the form of interpreted bytecode rather than machine-specific binaries to increase their interoperability between heterogeneous systems. A prominent example is Dalvik bytecode utilized by the Android system, which is currently the most popular platform for mobile and smart devices. This high-level representation of program semantics, however, suffers from the ease of reverse-engineering and tampering of apps when compared to native code. The ease of reverse-engineering is a main reason of the high rates of software piracy and malware distribution in the Android ecosystem.

The key goals of this subproject are the protection of intellectual property of software developers and the mitigation of malware threats. TP5 is a FORSEC interface project and as such a member of Cluster 2 as well as Cluster 3.

3.5.2 Results Achieved

The main contributions achieved within the scope of TP5 belong to the topics of malware detection and software protection for the Android mobile platform. Above that, results in the fields of obfuscation theory and mobile security in general have been achieved.

For malware detection and analysis, existing solutions were evaluated, identifying their weaknesses and room for improvements. Furthermore, new approaches were proposed. One subject of the examination, Google's VerifyApps also known as the package verification service, is an anti-malware tool targeting apps installed from sources alternative to the official Google Play store (Naumann et al. 2015). The practical evaluation of this software has indicated, and the reverse engineering of its internal functioning has confirmed, that it does not provide a sufficient protection as the detection process is based on an app's metadata only, excluding the app's code and behavior from the detection. Indeed, the detection rate even for known unmodified malware samples did not exceed 42%. Re-zipping the application package, which results only in the change of its hash signature but not its contents, reduced the detection rate to below 3%. Finally, application of static obfuscation techniques, such as identifier renaming, control and data flow obfuscation, and object-oriented design obfuscation, has completely prevented the detection of *all* 6,000 samples utilized in this test. Furthermore, the constant runtime monitoring of already installed apps, advertised by Google, was also found to be ineffective and partially dysfunctional.

Another evaluation of anti-malware solutions for Android, which also covered Bouncer used to sanitize the official Google Play store, has demonstrated a practical way of their circumvention. The analysis and detection tools were found to exhibit certain properties, which distinguish their environments from actual user devices (Maier et al. 2014, 2015). As a consequence, so called *split personality malware* can load and execute its malicious payloads only if executed on a real device, while behaving in a benign way inside the sandboxing analysis tools.

Furthermore, a comparative evaluation of attributes used for machine learning based detection of Android malware was performed, which covered both previously known and newly proposed attributes (Hahn et al. 2016). In the outcome of this evaluation, performed using 20,000 benign and malicious apps, the best performing single attribute set was found to be Android permissions, able to provide accuracy of 96%, although at cost of about 4% false positive rate. Among the most promising attribute combinations, the highest ranking was reached by permissions, intents, and app components and permissions together with opcode frequencies, demonstrating accuracy of about 97% and 96%, respectively and false positive rates of 2% and 1%, respectively.

For detection of Android malware a new approach based on software complexity metrics was proposed (Protsenko and Müller 2014). Software complexity metrics, known from the field of software engineering, serve as an indication of software implementation and design quality. In our approach, we utilize such metrics as McCabe's cyclomatic number, dependency degree, and the object-oriented design metrics suite by Chidamber and Kemerer as an attribute source for machine learning based detection. A practical evaluation of this technique, performed on a set of about 32,000 apps has confirmed its good detection performance, with a low false positive rate compared to permissions-based detection.

The practical evaluation of online dynamic analysis frameworks, also referred to as sandboxes, has considered MobileSandbox, Sanddroid, and Andrubis, for example (Busch et al. 2015). The three tools offer quite similar functionalities, such that no one could be distinguished as a clearly leading solution. As the main limitations, shared by all of the tools, we have identified their bounded abilities to process native code and readability of the generated reports.

Additionally, a dynamic analysis approach based on function hooking for Android's ART runtime was developed (Dresel et al. 2016). This technique inserts hooks into native code, thus being able to cover both native libraries and Java code AOT-compiled by ART. The hooking of a function is performed by replacing its first instruction with a breakpoint or illegal instruction, which on their execution cause a signal rise, thus delegating execution to the user-defined signal handler. The practical evaluation of this analysis approach has indicated its applicability to over 80% of Java functions and an acceptable performance penalty.

Further contributions include the introduction of novel software protection techniques for Android. These results are based on our early work, represented by the PANDORA randomized bytecode transformation framework (Protsenko and Müller 2013). This tool applies several bytecode diversifying transformations, modifying the control and data flow of functions and the object-oriented design in a random manner. Thus, it is possible to generate multiple versions of the same app with different code structure, which can be utilized to conceal plagiarized software or impede detection of malicious apps.

The newly proposed protection techniques capitalize on the advantages of native code with respect to software reverse engineering compared to bytecode, one reason for which is defined by the overall high-level structure of Dalvik bytecode, making it easily decompilable to the Java source code. Beyond that, the most available tools facilitating reverse engineering and analysis of Android apps, for example Harvester and TaintDroid, do not support processing of native code with the same degree of precision.

First, static obfuscation techniques based on native code were introduced (Protsenko and Müller 2015). Using four simple transformations, namely opaque predicates based on values received from native code, control flow flattening (partially outlined into the native library), and method call and field access indirection through native code, one can introduce artificial interdependencies between bytecode and the native library. Thus, the implemented transformations force the adversary to analyze or reverse engineer code for both execution domains in order to gain a complete understanding of the program's functioning, making precise inter- and intraprocedural control and data flow analysis based solely on bytecode infeasible.

Second, we make use of native code's capability to modify and inspect bytecode at runtime, provided by the fact that both execution environments share the same process (Protsenko et al. 2015). This allows for implementation of techniques providing strong protection against both dynamic and static analyses. For the purpose of tamperproofing, bytecode integrity checks based on checksum calculation can be integrated at any point of execution, preventing unwanted modification of app's behavior. Complementary, an obfuscation approach referred to as self-decrypting or self-modifying code can be enabled, decrypting the bytecode of functions prior to their execution and re-encrypting it after, thus making the cleartext bytecode of an app never completely exposed to the adversary at runtime. Based on two state-of-the-art taxonomies of dynamic obfuscation methods, our approach can be assigned the highest level of reached protection.

Third, an infrastructure for software protection was designed targeting the newer Android runtime ART, designated by the ahead-of-time (AOT) compilation of app's bytecode which usually happens at the installation time. The proposed infrastructure moves the AOT compilation from the untrusted user device to the trusted server operated by the app developer. This compilation process enables a novel obfuscation approach based on bytecode stripping: after being compiled, bytecode instructions can be removed from the app's executable file. Additionally, the described infrastructure makes possible the use of other protection techniques, such as self-modifying machine code, device-specific tamperproofing, and checksum-based integrity checks of code.

Aiming to gain a better understanding of the effect obfuscation has on the process of software reverse engineering, experiments have been performed with participants of various backgrounds and skills involved in solving program comprehension tasks. The results confirm a negative impact of the investigated obfuscation techniques, namely opaque predicates and name overloading, on the performance of humans in program reverse engineering (Zhuang et al. 2014).

Additionally, security evaluation of online banking based on so called mobile TAN approach performed on the example of Sparkasse app has demonstrated the conceptual weakness, which can allow the potential adversary controlling the user's device to redirect money transactions to his or her own account (Haupert and Müller 2016).

3.5.3 Contribution to FORSEC Research Alliance

Within the context of cluster STAR, focusing on smart environments, such as smart cities and smart homes, Android plays a crucial role as a platform operating not only the largest share of smartphones and tablet PCs, but also a variety of other smart devices, including refrigerators, television sets, and alarm clocks. TP5 contributes with identification of key security issues on Android, such as the widespread of malicious software and application piracy, evaluation of existing solutions, such as anti-malware and app analysis tools, discovery of their weaknesses and suggestion of improvements. And finally, new methodologies and techniques are designed and implemented, protecting the intellectual property of app developers and helping to mitigate the malware threat, in summary achieving the overall increase of security in the Android ecosystem. In this way, the numerous applications in the smart environments scenarios are facilitated and secured.

In the cluster CLOUD, addressing security in cloud environments, TP5 contributes with its knowledge in malware detection and analysis, gained in its main research topics concerning security of smartphones and Android in particular. Thus, in our joint publication together with other cluster members a novel approach to the mitigation of the malware threat in the cloud, referred to as CloudIdea (Fischer et al. 2015) was proposed, which makes use of lightweight monitoring techniques for attack detection enhanced with a capability of more sophisticated analysis techniques which can be applied after migrating the affected VM on a dedicated analysis host. Further joint research results in cluster CLOUD include a classification of malware and attacks on Infrastructure as a Service (IaaS) cloud environments with regard to the involved entities, relevance, and business impact.

3.5.4 Beyond FORSEC

Based on the previously outlined results already achieved in TP5, new research projects have been initiated; some of them are outlined next.

Considering the measures against software piracy on Android, aside of software protection techniques, one can address the question of a reliable identification of app repackaging. Currently, the state-of-the-art approaches either evaluate similarity of code or analyze the GUI structure using XML layout definition files. Since both are vulnerable against automated app transformations, which diversify an app's code or hide its resources from static analysis, we propose a dynamic approach which, while automatically executing an app, gathers the representation of its visual appearance as the values of the perceptual hash functions computed for the screenshots of an app. This technique is expected to be resilient against known automated app modifications.

Aiming to facilitate the process of app reverse engineering, e.g., for the purpose of malware analysis, we propose a framework for automated app deobfuscation, particularly targeting identifier renaming as the most widely used protection. The underlying approach is based on employing code similarity measurement to find unobfuscated of previously analyzed code similar or identical to the one being currently in focus of the analyst.

Other projects planned for future work include an empirical study of malicious native libraries on Android and their evolution, identification of private data leaks in Android apps by means of blackbox differential testing, and the visual representation of relationships between malware families and samples.

3.5.5 Publications

Baumann et al. (2017): Anti-ProGuard: Towards Automated Deobfuscation of Android Apps	
Abstract	A wide adoption of obfuscation techniques by Android application developers, and especially malware authors, introduces a high degree of complication into the process of reverse engineering, analysis, and security evaluation of third-party and potentially harmful apps. In this paper we present the early results of our research aiming to provide reliable means for automated deobfuscation of Android apps. According to the underlying approach, deobfuscation of a given app is performed by matching its code parts to the unobfuscated code stored in a database. For this purpose we apply well-known software similarity algorithms, such as SimHash and n-gram based ones. As a source of unobfuscated code can serve open source apps and libraries, as well as previously analyzed and manually deobfuscated code. Although the presented techniques are generic in their nature, our current prototype mainly targets Proguard, as one of the most widely used protection tools for Android performing primarily renaming obfuscation. The evaluation of the presented Anti-ProGuard tool witnesses its effectiveness for the considered task and supports the feasibility of the proposed approach.
Citation	Baumann, Richard; Protsenko, Mykola; Müller, Tilo (2017): Anti-ProGuard: Towards Automated Deobfuscation of Android Apps. In: *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems (SHCIS 2017).* Neuchâtel, Switzerland.

Busch et al. (2015): Automated Malware Analysis for Android: A Comparative Evaluation	
Abstract	In this paper, we show to what extent automatically generated reports for Android apps can help analyzing potentially malicious behavior. We generated reports using three well known analysis platforms for eleven malware and six goodware samples. Using the analysis reports generated by Andrubis, Mobile-Sandbox and SandDroid, we firstly evaluate each platform's ability to express information about an app's maliciousness. It turns out that no appropriate classification in goodware and malware can be performed by the assessed frameworks without relying on third party mostly signature based detection engines. Secondly, we discuss the contents presented within the generated malware reports and take them as a basis for comparing the frameworks. This comparison leads to the conclusion that among the assessed frameworks no truly superior solution exists.
Citation	Busch, M., Protsenko, M., and Müller, T. 2015. “Automated Malware Analysis for Android: A Comparative Evaluation,” in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015),* Vienna, Austria.

Dresel et al. (2016): ARTIST: The Android Runtime Instrumentation Toolkit	
Abstract	Smartphones are becoming more and more ubiquitous in the modern world, entrusted with such sensitive information as the user's location and banking data. Since Android is the most widespread smartphone platform, reliable and versatile means for Android application analysis are of great importance. Most of the existing code instrumentation approaches for Android suffer from two important shortcomings: the need for root access and limited support for the new Android Runtime (ART). We aim to fill this gap by proposing ARTIST, the Android Runtime Instrumentation Toolkit1. ARTIST is a framework that allows analysts to easily monitor the execution of Java and native code using native instrumentation techniques. ARTIST, to the best of our knowledge, is the first tool allowing monitoring of both native and Java code with the same instrumentation technique. ARTIST provides two methods to locate instrumentation targets. First, it can parse OAT executable files in memory to find classes and methods of interest. This allows monitoring a specific set of Java methods. Second, ARTIST can locate internal structures of the Android Runtime in memory. Monitoring function pointers found in these allows the user to track specific interactions of Java code with the Android Runtime. We evaluate the applicability of native instrumentation for Java code using a set of the most popular Android apps. The results show that over 80% of the tested Java methods are targetable using this approach. The performance impact, estimated with the CaffeineMark benchmark suite, does not exceed 20% and therefore can be considered generally acceptable.
Citation	Dresel, L., Protsenko, M., and Müller, T. 2016. “ARTIST: The Android Runtime Instrumentation Toolkit,” in *Proceedings of the 11th International Conference on Availability, Reliability and Security (ARES 2016),* Salzburg, Austria.
URL	http://ieeexplore.ieee.org/document/7784561/

Freiling et al. (2014): An Empirical Evaluation of Software Obfuscation Techniques Applied to Android APKs	
Abstract	We investigate the problem of creating complex software obfuscation for mobile applications. We construct complex software obfuscation from sequentially applying simple software obfuscation methods. We define several desirable and undesirable properties of such transformations, including idempotency and monotonicity. We empirically evaluate a set of 7 obfuscation methods on 240 Android Packages (APKs). We show that many obfuscation methods are idempotent or monotonous.
Citation	Freiling, F., Protsenko, M., and Zhuang, Y. 2014. “An Empirical Evaluation of Software Obfuscation Techniques applied to Android APKs,” in *Proceedings of the International Workshop on Data Protection in Mobile and Pervasive Computing (DAPRO 2014),* Beijing, China.
URL	https://link.springer.com/chapter/10.1007/978-3-319-23802-9_24

Freiling et al. (2017): Characterizing Loss of Digital Evidence Due to Abstraction Layers	
Abstract	We study the problem of evidence collection in environments where abstraction layers are used to organize data storage. Based on a formal model, the problem of evidence collection is defined as the task to reconstruct high-level from low-level storage. We investigate the conditions under which different levels of evidence collection can be performed and show that abstraction layers, in general, make it harder to acquire evidence. We illustrate our findings by describing several practical scenarios from file systems, memory management, and disk volume management.
Citation	Freiling, F., Glanzmann, T., and Reiser, H. P. 2017. “Characterizing loss of digital evidence due to abstraction layers,” Digital Investigation (20), pp. 107–115.
URL	http://www.sciencedirect.com/science/article/pii/S1742287617300427

Groß and Müller (2017): Protecting JavaScript Apps from Code Analysis	
Abstract	Apps written in JavaScript are an easy target for reverse engineering attacks, e.g. to steal the intellectual property or to create a clone of an app. Unprotected JavaScript apps even contain high level information such as developer comments, if those were not explicitly stripped. This fact becomes more and more important with the increasing popularity of JavaScript as language of choice for both web development and hybrid mobile apps. In this paper, we present a novel JavaScript obfuscator based on the Google Closure Compiler, which transforms readable JavaScript source code into a representation much harder to analyze for adversaries. We evaluate this obfuscator regarding its performance impact and its semantics-preserving property.
Citation	Groß, Tobias; Müller, Tilo (2017): Protecting JavaScript Apps from Code Analysis. In: *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems (SHCIS 2017)*. Neuchâtel, Switzerland.

Hahn et al. (2016): Comparative Evaluation of Machine Learning-based Malware Detection on Android	
Abstract	The Android platform is known as the market leader for mobile devices, but it also has gained much attention among malware authors in recent years. The widespread of malware, a consequence of its popularity and the design features of the Android ecosystem, constitutes a major security threat currently targeted by the research community. Among all counter methods proposed in previous publications, many rely on machine learning algorithms based on statically extracted attributes from an app. Machine learning, which is also inspired by the developed field of desktop malware detection, has proven to be a promising approach for fighting Android malware. Many publications, however, rely on different data sets for different application attributes, rendering the comparison of them difficult. Furthermore, there exist attribute sets known from the desktop world which have not been ported to Android yet. In this paper, we aim to step towards filling this gap by assessing the effectiveness of the total number of 11 attribute sets, including those never evaluated on Android before, using a consistent data set of 10,000 apps. Our comparative evaluation provides a ranking for the single attribute sets according the detection performance they can reach, and suggests the most effective combination of all attributes.
Citation	Hahn, S., Protsenko, M., and Müller, T. 2016. “Comparative Evaluation of Machine Learning-based Malware Detection on Android,” in Sicherheit 2016.
URL	https://subs.emis.de/LNI/Proceedings/Proceedings256/79.pdf

Maier et al. (2014): Divide-and-Conquer: Why Android Malware Cannot Be Stopped	
Abstract	In this paper, we demonstrate that Android malware can bypass all automated analysis systems, including AV solutions, mobile sandboxes, and the Google Bouncer. We propose a tool called Sand-Finger for the fingerprinting of Android-based analysis systems. By analyzing the fingerprints of ten unique analysis environments from different vendors, we were able to find characteristics in which all tested environments differ from actual hardware. Depending on the availability of an analysis system, malware can either behave benignly or load malicious code at runtime. We classify this group of malware as Divide-and-Conquer attacks that are efficiently obfuscated by a combination of fingerprinting and dynamic code loading. In this group, we aggregate attacks that work against dynamic as well as static analysis. To demonstrate our approach, we create proof-of-concept malware that surpasses up-to-date malware scanners for Android. We also prove that known malware samples can enter the Google Play Store by modifying them only slightly. Due to Android's lack of an API for malware scanning at runtime, it is impossible for AV solutions to secure Android devices against these attacks.
Citation	Maier, D., Müller, T., and Protsenko, M. 2014. “Divide-and-Conquer: Why Android Malware cannot be stopped,” in *Proceedings of the 9th International Conference on Availability, Reliability and Security (ARES 2014)*, Fribourg, Switzerland.
URL	http://ieeexplore.ieee.org/document/6980261/

Maier et al. (2015): A Game of Droid and Mouse: The Threat of Split-Personality Malware on Android	
Abstract	In the work at hand, we first demonstrate that Android malware can bypass current automated analysis systems, including AV solutions, mobile sandboxes, and the Google Bouncer. A tool called Sand-Finger allowed us to fingerprint Android-based analysis systems. By analyzing the fingerprints of ten unique analysis environments from different vendors, we were able to find characteristics in which all tested environments differ from actual hardware. Depending on the availability of an analysis system, malware can either behave benignly or load malicious code dynamically at runtime. We also have investigated the widespread of dynamic code loading among benign and malicious apps, and found that malicious apps make use of this technique more often. About one third out of 14,885 malware samples we analyzed was found to dynamically load and execute code. To hide malicious code from analysis, it can be loaded from encrypted assets or via network connections. As we show, however, even dynamic scripts which call existing functions enable an attacker to execute arbitrary code. To demonstrate the effectiveness of both dynamic code and script loading, we create proof-of-concept malware that surpasses up-to-date malware scanners for Android and show that known samples can enter the Google Play Store by modifying them only slightly.
Citation	Maier, D., Protsenko, M., and Müller, T. 2015. “A Game of Droid and Mouse: The Threat of Split-Personality Malware on Android,” in Computers & Security (COSE).
URL	http://www.sciencedirect.com/science/article/pii/S0167404815000656

Naumann et al. (2015): Google Verify Apps: The Illusion of Security?	
Abstract	In this paper we analyze Verify Apps, which is the standard anti-virus software for Android offered by Google. Verify Apps should protect the user from malicious apps which are installed from other sources than Google's Play Store. To get more information about the internals of Verify Apps, we tested it with 6103 malware apps, each modified with four different obfuscation techniques. In addition, we examined its functionality in detail by reverse engineering. Verify Apps recognized about 42 percent of the original malware samples, but even a such simple transformation as re-zipping the app, which only affects the hash signatures of a file, resulted in detection rate dropping to less than 3 percent. After the application of static obfuscation techniques, none of the samples could be detected. Moreover, we experienced practical problems with Verify Apps: It stops working after identifying eleven apps as malware, and the verification of already installed apps, which is one of its new features according to Google, could not be observed.
Citation	Naumann, J., Protsenko, M., and Müller, T. 2015. “Google Verify Apps: The Illusion of Security?,” in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015),* Vienna, Austria.

Protsenko and Müller (2013): PANDORA Applies Non-Deterministic Obfuscation Randomly to Android	
Abstract	Android, a Linux-based operating system, is currently the most popular platform for mobile devices like smart-phones and tablets. Recently, two closely related security threats have become a major concern of the research community: software piracy and malware. This paper studies the capabilities of code obfuscation for the purposes of plagiarized software and malware diversification. Within the scope of this work, the PANDORA (PANDORA Applies Non-Deterministic Obfuscation Randomly to Android) transformation system for Android bytecode was designed and implemented, combining techniques for data and object-oriented design obfuscation. Our evaluation results indicate deficiencies of the malware detection engines currently used in 46 popular antivirus products, which in most cases were not able to detect samples obfuscated with PANDORA. Furthermore, this paper reveals shortcomings of the Androsim tool and potentially other static software similarity algorithms, recently proposed to address the piracy problem in Android.
Citation	Protsenko, M., and Müller, T. 2013. “PANDORA Applies Non-Deterministic Obfuscation Randomly to Android,” in *Proceedings of the 8th International Conference on Malicious and Unwanted Software (Malware 2013),* Fajardo, Puerteo Rico.
URL	http://ieeexplore.ieee.org/document/6703686/

Protsenko and Müller (2014): Android Malware Detection based on Software Complexity Metrics	
Abstract	In this paper, we propose a new approach for the static detection of Android malware by means of machine learning that is based on software complexity metrics, such as McCabe's Cyclomatic Complexity and the Chidamber and Kemerer Metrics Suite. The practical evaluation of our approach, involving 20,703 benign and 11,444 malicious apps, witnesses a high classification quality of our proposed method, and we assess its resilience against common obfuscation transformations. With respect to our large-scale test set of more than 32,000 apps, we show a true positive rate of up to 93% and a false positive rate of 0.5% for unobfuscated malware samples. For obfuscated malware samples, however, we register a significant drop of the true positive rate, whereas permission-based classification schemes are immune against such program transformations. According to these results, we advocate for our new method to be a useful detector for samples within a malware family sharing functionality and source code. Our approach is more conservative than permission-based classifications, and might hence be more suitable for an automated weighting of Android apps, e.g., by the Google Bouncer.
Citation	Protsenko, M., and Müller, T. 2014. “Android Malware Detection based on Software Complexity Metrics,” in *Proceedings of the 11th International Conference on Trust, Privacy & Security in Digital Business (TrustBus 2014),* Munich, Germany.
URL	https://link.springer.com/chapter/10.1007/978-3-319-09770-1_3

Protsenko et al. (2015): Dynamic Self-Protection and Tamperproofing for Android Apps using Native Code	
Abstract	With over one billion sold devices, representing 80% market share, Android remains the most popular platform for mobile devices. Application piracy on this platform is a major concern and a cause of significant losses: about 97% of the top 100 paid apps were found to be hacked in terms of repackaging or the distribution of clones. Therefore new and stronger methods aiming to increase the burden on reverse engineering and modification of proprietary mobile software are required. In this paper, we propose an application of the Android native code component to implement strong software self-protection for apps. Within this scope, we present three dynamic obfuscation techniques, namely dynamic code loading, dynamic re-encryption, and tamper proofing. We provide a practical evaluation of this approach, assessing both the cost and efficiency of its achieved protection level. Our results indicate that with the proposed methods one can reach significant complication of the reverse-engineering process, while being affordable in terms of execution time and application size.
Citation	Protsenko, M., Kreuter, S., and Müller, T. 2015. “Dynamic Self-Protection and Tamperproofing for Android Apps using Native Code,” in *Proceedings of the 10th International Conference on Availability, Reliability and Security (ARES 2015)*, Toulouse, France.
URL	http://ieeexplore.ieee.org/document/7299906/

Protsenko and Müller (2015): Protecting Android Apps against Reverse Engineering by the Use of the Native Code	
Abstract	Having about 80 % of the market share, Android is currently the clearly dominating platform for mobile devices. Application theft and repackaging remains a major threat and a cause of significant losses, affecting as much as 97 % of popular paid apps. The ease of decompilation and reverse engineering of high-level bytecode, in contrast to native binary code, is considered one of the main reasons for the high piracy rate. In this paper, we address this problem by proposing four static obfuscation techniques: native opaque predicates, native control flow flattening, native function indirection, and native field access indirection. These techniques provide a simple and yet effective way of reducing the task of bytecode reverse engineering to the much harder task of reverse engineering native code. For this purpose, native function calls are injected into an app's bytecode, introducing artificial dependencies between the two execution domains. The adversary is forced to analyze the native code in order to be able to comprehend the overall app's functionality and to successfully launch static and dynamic analyses. Our evaluation results of the proposed protection methods witness an acceptable cost in terms of execution time and application size, while significantly complicating the reverse-engineering process.
Citation	Protsenko, M., and Müller, T. 2015. "Protecting Android Apps against Reverse Engineering by the Use of the Native Code," in *Proceedings of the 12th International Conference on Trust, Privacy & Security in Digital Business (TrustBus 2015),* Valencia, Spain.
URL	https://link.springer.com/chapter/10.1007/978-3-319-22906-5_8

Zhuang and Freiling (2015): Approximating Optimal Software Obfuscation for Android Applications	
Abstract	In the context of software protection, we study the problem of automatically obfuscating a given program to a given target level of “difficulty”. We measure difficulty by utilizing software complexity metrics. We formalize the search problem and argue that current informed search algorithms cannot be used for our purpose, because the number of evaluated search candidates should be minimal and their actual complexities cannot be predicted with certainty. Within a framework for program obfuscation for Android APKs, we empirically evaluate two different algorithms that search for an obfuscated version satisfying a conjunction of target complexity metrics. We show that a first algorithm whose predictions rely on mean values is outperformed by a second algorithm based on Bayes theorem.
Citation	Zhuang, Y., and Freiling, F. 2015. “Approximating Optimal Software Obfuscation for Android Applications,” in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015),* Vienna, Austria.

Zhuang et al. (2014): An(other) Exercise in Measuring the Strength of Source Code Obfuscation	
Abstract	We experimentally compare the strength of different source code obfuscation techniques by measuring the performance of human analysts. We describe an experimental setup by which it is possible to compare different obfuscation techniques with each other. As techniques, we considered name overloading and opaque predicates, as well as the combination of both. While the results are interesting and show that increased levels of obfuscation decrease the performance of humans, only one result (the use of name overloading) was statistically significant.
Citation	Zhuang, Yan; Protsenko, Mykolai, Müller, Tilo, and Freiling, Felix (2014): An(other) Exercise in Measuring the Strength of Source Code Obfuscation. In: *Proceedings of the 1st Workshop on Security in highly connected IT systems (SHCIS 2014)*, Munich, Germany.
URL	http://ieeexplore.ieee.org/document/6974868/

3.6 TP6 – Security Awareness

3.6.1 Project Overview

User behavior greatly influences the effectiveness of IT security. For example, when users grant permissions to smartphone apps, this could lead to data leaks. The problem is especially severe when the smartphone is both used in private and work context. Another example is phishing, where attack success directly depends on user interaction. Users should be aware of the dangers and consequences of their actions in IT environments. The first TP goal was therefore to understand the meaning of security awareness for the users. Further goals were to investigate the current security and privacy awareness of users and develop methods for its improvement. To this end, we first conducted a literature review on existing methods of raising security awareness. We then researched the existing awareness of users, especially with respect to the usage of social media and smartphones. For the former, we investigated susceptibility of social media users to phishing. For the latter, we compared the privacy and security awareness of users of the two most commonly used operating systems, Android and iOS. Here, not only users at home were studied, but also how smartphone usage is handled in companies with focus on security-relevant topics. The latter topic was studied in cooperation with TP10. Furthermore, the usability and user acceptance of existing security- and privacy-enhancing mechanisms was studied. Here, we considered smartphone permissions as well as an advanced Privacy Enhancing Technology (PET) called "attribute-based credentials". TP6 also focused on new methods for raising awareness. We helped TP8 in the development and evaluation of new user interfaces that support users in assessment of seller trustworthiness on eBay. Moreover, in cooperation with TP5 we conducted a user study on the effectiveness of software obfuscation methods.

3.6.2 Results Achieved

We first conducted a literature review with the goals of systematization of existing security awareness approaches (Hänsch and Benenson 2014). For this project, we analyzed how security awareness it treated in the field of computer science and what implications this has on security awareness training methods. The result of our analysis is that in general, three types of security awareness are discussed in literature. The first group of papers puts its emphasis on teaching users how to recognize a threat. Some of them even explicitly argue that more is not needed to be "aware". The second group enriches this understanding of awareness by adding the need to know solutions to these threats. The third group of publications argues that not only knowing but also acting according to these solutions is necessary in order to be "aware". Interestingly, all analyzed papers used or suggested different types of evaluations of security awareness measures. This led us to the result that the term "security awareness" is not clearly defined yet. Organizations interested in initiating an awareness program should therefore clearly communicate which type of awareness they aim for to get an individually designed program for their needs. Furthermore, we analyzed capacities of the users for successful execution of security tasks (Benenson, Lenzini, et al. 2015). The evidence from the literature suggests that even security aware users may not be able to execute security tasks due to the cognitive and physical limits of humans. We propose a conceptual framework for evaluation of human capacities in security that also assigns systems to complexity categories according to their security and usability.

We also considered security and privacy awareness in social media (publications (Benenson, Gassmann, et al. 2017; Benenson, Girard, Hintz, et al. 2014; Hintz et al. 2014) and student theses), especially focusing on phishing susceptibility of Facebook users compared to users of traditional email communication. To this end, we conducted two experiments where we sent to over 1700 university students an email or a Facebook message with a link from a non-existing person, claiming that the link leads to the pictures from the party last week. When clicked, the corresponding webpage showed the "access denied" message. We registered the click rates, and later sent a questionnaire to the participants a questionnaire that first assessed their security awareness, and then informed them about the experiment and asked them about the reasons for their clicking behavior. When addressed by first name, 56% of email and 38% of Facebook recipients clicked. When not addressed by first name, 20% of email and 42% of Facebook recipients clicked. By far the most frequent reason for clicking was curiosity about the content of the pictures or the personality of the sender, followed by the explanations that the content or context of the message fits the current life situation of the person, such as actually knowing somebody with this name, or having been at a party with unknown people last week. Although these studies revealed susceptibilities to scam in some people, and the reasons behind their susceptibility, but we think that the lessons learned are broader. By a careful design and timing of a message, it should be possible to make virtually any person click on a link, as any person will be curious about something, or interested in some topic, or be in a life situation that fits the message's content and context. In the long run, relying on technical in-depth defense may be a better solution, and more research and evidence is needed to determine which level of defense non-expert users are able to achieve through security education and training.

Considering privacy and security awareness of smartphone users (publications (Reinfelder and Weishäupl 2016; Reinfelder et al. 2014; Russ et al. 2017) and student theses), we investigated smartphone usage in private as well as in the organizational context. A survey with 700 Android and iPhone users compared their security and privacy awareness when handling apps (Reinfelder et al. 2014). We found out that Android users seem to be more aware of the risks associated with the app usage than iPhone users. For example, iPhone users almost never consider the possibility of apps sending premium-rate SMS or causing other hidden costs. Furthermore, Android users more often mention security, trust and privacy issues as important factors when they decide to use a new app. We hypothesize that the cause of these differences is likely to arise through differences in app market policies, in app review processes and in presentation of data usage by the apps. On the other hand, Android app permissions have been widely criticized for their poor usability, whereas the runtime permissions model of iOS received relatively low attention in the usable security community. Since October 2015, Android also implements the runtime permissions model. To investigate how this change is perceived by the Android users, we conducted a survey with over 800 respondents (Russ et al. 2017). We compare perception and reported usage of the respective permissions models by three groups: users of old Android, runtime Android and iOS permissions. The results indicate that both permissions types are utilized by the users for decision making regarding app installations and usage. However, runtime permissions in Android and iOS are perceived as more useful than the old Android permissions. The users also show a more positive emotional attitude to the runtime permissions model.

In the organizational context, the Dynamic Security Success Model for smartphone security developed was developed in cooperation with TP10 (Reinfelder and Weishäupl 2016). This model is now being evaluated in various companies. To this end, we conducted interviews with security professionals in seven companies about how they manage smartphone security in their organizations, and what is the role of the users in this process. The results show that security professionals consider users as unwilling to protect company from the attacks, and feel that security is often sacrificed for the sake of business needs. All but one companies lack communication channels to the end users, especially considering how to receive feedback from them. We conclude that improved communication between security professionals and end users is needed to raise security awareness and achieve better security in companies.

Further a user acceptance study of a system for anonymous credentials was conducted in cooperation with the EU project ABC4Trust (Benenson, Girard, et al. 2015; Benenson, Girard, Krontiris, et al. 2014). We developed a user acceptance model for anonymous credentials based on the Technology Acceptance Model (TAM). We introduce five new constructs into the TAM: Perceived Usefulness for the Primary Task (PU1), Perceived Usefulness for the Secondary Task (PU2), Situation Awareness, Perceived Anonymity and Understanding. We conduct an evaluation of our model in the concrete scenario of a university course evaluation with 30 students. We conclude that PU1 is the most important factor of user adoption, outweighing the usability and the usefulness of the deployed technology (PU2). Moreover, correct Understanding of the underlying technologies seems to play a less important role than a user interface of the system that clearly conveys to the user which data are transmitted when and to which party (Situation Awareness).

In (Sänger et al. 2016) we have shown that users can detect fraudulent sellers on eBay better when they use a new visualization for the seller reputation developed by TP8 compared to the standard eBay visualization. The study was conducted in cooperation with the University College London with 40 UK and 41 German participants. Half of them used reputation visualization similar to the one that is now being used on eBay. The rest worked with the new interface developed by TP8. Each participant had to buy four products. For each product, the participants could choose between two sellers. Both of them had a similar reputation score, although for each case one of the sellers was involved in a specific kind of fraud: building up good reputation with cheap products while cheating on expensive ones, cheating only when selling a certain type of product or receiving only negative feedback recently. In the fourth case one of the sellers have never sold a product of that type, which adds uncertainty to the possible purchase, while the other one already have and received positive feedback for it. Our results showed that the overall detection of malicious sellers was 39 % for participants using the old system compared to 65 % using the new visualization. Especially users with low experience further benefited from using this new interface, as their detection of malicious sellers was 63 % compared to 28 % for those who used the old system. Further, the new interface was well accepted by those who used it. Most participants agreed that they would like to use the new system if it was available on eBay, as it helps them to get a better insight into sellers' trustworthiness. They also think the benefits of using it are bigger than the effort to do so.

In cooperation with TP5 we studied how adding software obfuscation on source code effects the attackers' ability to understand it (Benenson, Freiling, et al. 2017). In this user study, participants were given two programs to analyze. For each program two tasks, which indicate the participants understanding of the code, had to be solved. The first program was given in clear code, while the second one was obfuscated either with the "name obfuscation" or "opaque predicates" technique. Our results show as expected that experienced users are better at understanding how clear code programs work compared to our participants with less experience. This gap although shrinks when one of the obfuscation techniques is applied. We further analyzed which tools are being used by our participants in order to understand the code they are analyzing. Our results suggest that participants with higher knowledge spent more time debugging in general. They also initiate the debugging process more often than our participants with lower experience. Considering obfuscation in general, if one of the obfuscation types was added, our participants needed more time and needed the help of features of the programming environment more often instead of relying on reading the code only. Concerning the differences between the type of obfuscation, we found out that our participants needed to switch between different source code files significantly more often when it was obfuscated with "name overloading" than the "opaque predicates" method. Overall our results implicate that the process of analyzing obfuscated software differs from the traditional understanding of clear code.

3.6.3 Contribution to FORSEC Research Alliance

In cooperation with TP8 we evaluated the usability of a new visualization method for eBay seller profiles (Sänger et al. 2016). We first conducted usability tests with a group of potential users, which helped TP8 to improve the user interface according to the received feedback. Furthermore, both TPs developed a user study to evaluate the effectiveness of the new system. To this end, TP8 worked on the technical aspects of the system, while TP6 developed the study design. Both TPs then conducted the usability study at FAU and at UCL (UK) with overall 81 participants.

Further, we helped to create questionnaire for a user study on a smart transportation smartphone app for smart cities as a cooperation project with TP2 and EU project RERUM. First, the specifications of the smart city project were analyzed, in order to find potential problems users could be confronted with or questions that could arise when using the system. Based on our analysis of the project specifications, and discussions with TP2 and RERUM, a questionnaire was developed and later used by RERUM for project evaluation.

In cooperation with TP5 we conducted a user study to find out how applying different obfuscation methods on a program changes the behavior of potential attackers (Benenson, Freiling, et al. 2017). TP5 delivered technical system to be used, while TP6 developed research questions and hypotheses for the study and worked on the study design and evaluation. Both TPs cooperatively the user study with over 70 participants and analyzed its results. TP5 concentrated on analysis of technical features, whereas TP6 focused on analysis of human behavior.

A systematic literature review of approaches to the management of smartphone security in companies was conducted in cooperation with TP10 and a theoretical framework for smartphone security in organizations was developed (Reinfelder and Weishäupl 2016). The resulting Dynamic Security Success Model is a combination of the DeLone & McLean Information Systems Success Model and Argyris' Organizational Learning Theory. This theoretical foundation combines the individual and the organizational impact of smartphone security measures with the learning perspective, allowing a company to respond to the ever changing security requirements of smartphones in organizations.

3.6.4 Beyond FORSEC

As part of (Sänger et al. 2016) we also conducted interviews with our participants to find out what they pay attention to when they buy things on eBay, and especially how they try to avoid fraud. By using Qualitative Content Analysis followed up by a hermeneutical interpretation we identified how buyers try to establish trustworthiness of a seller and estimate overall success of a transaction. An article about these results is in preparation, and a follow-up study is planned. The Dynamic Security Success Model developed in (Reinfelder and Weishäupl 2016) is now being evaluated in various companies. To this end, we conducted interviews with security professionals in seven companies about how they manage smartphone security in their organizations, and what is the role of the users in this process. These results are now in preparation to publication. Further evaluations with smartphone users in companies are being planned. Combining our results on security awareness (Hänsch and Benenson 2014) and phishing (Benenson, Gassmann, et al. 2017; Benenson, Girard, Hintz, et al. 2014; Hintz et al. 2014), we are planning to investigate the effectiveness of anti-phishing mechanisms in companies and develop improvements to the current state of awareness measures. By conducting interviews with security professionals and general users about their experience with protecting companies from the phishing threat will be conducted. Additionally, new metrics for quantitative assessment of the effects of the anti-phishing measures will be developed and evaluated. Moreover, further work on human factors in secure programming is to follow up the results of the software obfuscation experiment (Benenson, Freiling, et al. 2017). To further develop the privacy awareness topic, evaluation of user acceptance of privacy-enhanced cloud technologies is being conducted as a part of the Marie Skłodowska-Curie Innovative Training Network "Privacy and Usability" (EU Horizon 2020).

3.6.5 Publications

Benenson, Girard, Hintz, et al. (2014): Susceptibility to URL-based Internet attacks. Facebook vs. Email	
Abstract	The usage of social networking sites has been steadily increasing in the last decade. Communication via social networks has replaced email as the traditional means of electronic communication in many contexts. Accordingly, many types of Internet fraud also spread to social networks. In this work, we make the first to our knowledge direct comparison of users' susceptibility to attacks that involve clicking on dangerous links in Facebook messages versus in emails. We conducted a between-subjects quasi-experiment with 398 users where the users received links from strangers in a Facebook message or via email. We observed the respective clicking behavior and investigated users' attitudes to URL-based attacks by means of a post-experimental survey. Our results show that the communication medium (Facebook vs. email) leads to significant differences in attack susceptibility. Quite surprisingly, the success rate of email-based attacks is significantly higher.
Citation	Benenson, Z., Girard, A., Hintz, N., and Luder, A. 2014. “Susceptibility to URL-based Internet Attacks: Facebook vs. email,” in *Proceedings of the 6th IEEE International Workshop on SEcurity and SOCial Networking (SESOC 2014)*, Budapest, Hungary.
URL	http://ieeexplore.ieee.org/document/6815275/

Benenson, Girard, Krontiris, et al. (2014): User Acceptance of Privacy-ABCs: An Exploratory Study	
Abstract	In this work, we present the first statistical results on users' understanding, usage and acceptance of a privacy-enhancing technology (PET) that is called “attribute-based credentials”, or Privacy-ABCs. We identify some shortcomings of the previous technology acceptance models when they are applied to PETs. Especially the fact that privacy-enhancing technologies usually assist both, the primary and the secondary goals of the users, was not addressed before. We present some interesting relationships between the acceptance factors. For example, understanding of the Privacy-ABC technology is correlated to the perceived usefulness of Privacy-ABCs. Moreover, perceived ease of use is correlated to the intention to use the technology. This confirms the conventional wisdom that understanding and usability of technology play important roles in the user adoption of PETs.
Citation	Benenson, Z., Girard, A., Krontiris, I., Liagkou, V., Rannenberg, V., and Stamatiou, Y. 2014. “User acceptance of Privacy-ABCs: An Exploratory Study,” in *Proceedings of the 2nd International Conference on Human Aspects of Information Security, Privacy, and Trust*, Crete, Greece.
URL	https://link.springer.com/chapter/10.1007/978-3-319-07620-1_33

Benenson, Lenzini, et al. (2015): Maybe Poor Johnny Really Cannot Encrypt - The Case for a Complexity Theory for Usable Security	
Abstract	Psychology and neuroscience literature shows the existance of upper bounds on the human capacity for executing cognitive tasks and for information processing. These bounds are where, demonstrably, people start experiencing cognitive strain and consequently committing errors in the tasks execution. We argue that the usable security discipline should scientifically understand such bounds in order to have realistic expectations about what people can or cannot attain when coping with security tasks. This may shed light on whether Johnny will be ever be able to encrypt. We propose a conceptual framework for evaluation of human capacities in security that also assigns systems to complexity categories according to their security and usability. From what we have initiated in this paper, we ultimately aim at providing designers of security mechanisms and policies with the ability to say: "This feature of the security mechanism X or this security policy element Y is inappropriate, because this evidence shows that it is beyond the capacity of its target community".
Citation	Benenson, Z., Lenzini, G., Oliveira, D., Parkin, S., and Uebelacker, S. 2015. “Maybe Poor Johnny Really Cannot Encrypt - The Case for a Complexity Theory for Usable Security,” in New Security Paradigms Workshop, Twente, Netherlands.
URL	https://dl.acm.org/citation.cfm?id=2841120

Benenson, Gassmann, et al. (2017): Unpacking Spear Phishing Susceptibility	
Abstract	We report the results of a field experiment where we sent to over 1200 university students an email or a Facebook message with a link to (non-existing) party pictures from a non-existing person, and later asked them about the reasons for their link clicking behavior. We registered a significant difference in clicking rates: 20% of email versus 42.5% of Facebook recipients clicked. The most frequently reported reason for clicking was curiosity (34%), followed by the explanations that the message fit recipient's expectations (27%). Moreover, 16% thought that they might know the sender. These results show that people's decisional heuristics are relatively easy to misuse in a targeted attack, making defense especially challenging.
Citation	Benenson, Z., Gassmann, F., and Landwirth, R. 2017. "Unpacking Spear Phishing Susceptibility," in *Proceedings of the 21st International Conference on Financial Cryptography and Data Security*, Malta.

Benenson, Girard, et al. (2015): User Acceptance **Factors for Anonymous Credentials: An Empirical Investigation**	
Abstract	We describe theoretical development of a user acceptance model for anonymous credentials and its evaluation in a real-world trial. Although anonymous credentials and other advanced privacy-enhanced technologies (PETs) reached technical maturity, they are not widely adopted so far, such that understanding user adoption factors is one of the most important goals on the way to better privacy management with the help of PETs. Our model integrates the Technology Acceptance Model (TAM) with the considerations that are specific for security- and privacy-enhancing technologies, in particular, with their "secondary goal" property that means that these technologies are expected to work in the background, facilitating the execution of users' primary, functional goals. We introduce five new constructs into the TAM: Perceived Usefulness for the Primary Task (PU1), Perceived Usefulness for the Secondary Task (PU2), Situation Awareness, Perceived Anonymity and Understanding of the PET. We conduct an evaluation of our model in the concrete scenario of a university course evaluation. Although the sample size (30 participants) is prohibitively small for deeper statistical analysis such as multiple regressions or structural equation modeling, we are still able to derive useful conclusions from the correlation analysis of the constructs in our model. Especially, PU1 is the most important factor of user adoption, outweighing the usability and the usefulness of the deployed PET (PU2). Moreover, correct Understanding of the underlying PET seems to play a much less important role than a user interface of the system that clearly conveys to the user which data are transmitted when and to which party (Situation Awareness).
Citation	Benenson, Z., Girard, A., and Krontiris, I. 2015. "User Acceptance Factors for Anonymous Credentials: An Empirical Investigation," in *Proceedings of the 14th Workshop on the Economics of Information Security (WEIS 2015)*, Delft, Netherlands.

Hänsch and Benenson (2014): Specifying IT Security Awareness	
Abstract	IT users are faced with various threats on a daily basis. Unfortunately, not all possible dangers are known to them, such that the users fall an easy victim to attacks. For this reason, IT specialists demand for higher IT security awareness. Although researchers and practitioners exercise ongoing efforts in this area, their work often lacks a concise definition of the term "security awareness". Since there is no agreement on the term, different (and sometimes not compatible) ways of raising and measuring security awareness exist. This paper is an effort to give an overview of this phenomenon and to show that awareness in IT security is not standardized yet.
Citation	Hänsch, N., and Benenson, Z. 2014. “Specifying IT Security Awareness,” in *Proceedings of the 1st Workshop on Security in Highly Connected IT systems (SHCIS 2014)*, Munich, Germany.
URL	http://ieeexplore.ieee.org/document/6974870/

Krontiris et al. (2015): Privacy-ABCs as a Case for Studying the Adoption of PETs by Users and Service Providers	
Abstract	Although in the last years there has been a growing amount of research in the field of privacy-enhancing technologies (PETs), they are not yet widely adopted in practice. In this paper we discuss the socioeconomical aspects of how users and service providers make decisions about adopting PETs. The analysis is based on our experiences from the deployment of Privacy-respecting Attribute-based Credentials (Privacy-ABCs) in a real-world scenario. In particular, we consider the factors that affect the adoption of Privacy-ABCs as well as the cost and benefit trade-offs involved in their deployment and usage, as perceived by both parties.
Citation	Krontiris, I., Benenson, Z., Gerard, A., Sabouri, A., Rannenberg, K., and Schoo, P. 2015. “Privacy-ABCs as a Case for Studying the Adoption of PETs by Users and Service Providers,” in Annual Privacy Forum.
URL	https://link.springer.com/chapter/10.1007/978-3-319-31456-3_6

Reinfelder et al. (2014): Differences between Android and iPhone Users in Their Security and Privacy Awareness	
Abstract	This work compares Android and iPhone users according to their security and privacy awareness when handling apps. Based on an online survey conducted with over 700 German respondents (mostly university students) we found out that Android users seem to be more aware of the risks associated with the app usage than iPhone users. For example, iPhone users almost never consider the possibility of apps sending premium-rate SMS or causing other hidden costs. Furthermore, Android users more often mention security, trust and privacy issues as important factors when they decide to use a new app. We hypothesize that the cause of these differences they are likely to arise through differences in app market policies, in app review processes and in presentation of data usage by the apps.
Citation	Reinfelder, L., Benenson, Z., and Gassmann, F. 2014. “Differences between Android and iPhone Users in Their Security and Privacy Awareness,” in *Proceedings of the 11th International Conference on Trust, Privacy & Security in Digital Business (TrustBus 2014)*, Munich, Germany.
URL	https://link.springer.com/chapter/10.1007/978-3-319-09770-1_14

Reinfelder and Weishäupl (2016): A Literature Review on Smartphone Security in Organizations Using a New Theoretical Model-The Dynamic Security Success Model	
Abstract	Smartphones have become an important part of organizational IT infrastructures including benefits such as increased productivity as well as IT security risks. These risks are mainly related to unauthorized access to corporate data. Integrating smartphones in organizations regarding security involves a sequence of decisions, ranging from the integration approach (smartphones owned by employees or by the organization) to specific security measures implemented on the devices. This is an ongoing process making constant adaption necessary due to progressive development of hard- and software and due to new security risks arising. We propose the Dynamic Security Success Model (DSSM) – a combination of the DeLone & McLean Information Systems Success Model and Argyris' Organizational Learning Theory. This theoretical foundation combines the individual and the organizational impact of smartphone security measures with the learning perspective, allowing a company to respond to the ever changing security requirements of smartphones in organizations. Based on the DSSM, existing literature is reviewed and research gaps are derived for future work.
Citation	Reinfelder, L., and Weishäupl, E. 2016. "A Literature Review on Smartphone Security in Organizations using a new theoretical Model-The Dynamic Security Success Model," in *Proceedings of the 20th Pacific Asia Conference on Information Systems (PACIS 2016),* Chiayi, Taiwan.
URL	http://aisel.aisnet.org/pacis2016/59/

3.7 TP7 – Identity 3.0

3.7.1 Project Overview

The project deals with new approaches for in-house Identity and Access Management (IAM). It has the goal to contribute to more efficient management of IAM structures. Thus, it acts against currently existing identity and access privilege chaos. Identity 3.0 aims at developing measures and methods in order to help organizations to securely and efficiently manage private and corporate identities.

As such, among other intermediate results, the project developed solutions for various coupled problems within IAM, aiming at providing a toolset for existing IAM challenges. Among the twelve published results, four major contributions have been developed within the project: An approach for role quality assessment, for role optimization and for policy optimization as well as a migration guide for aiding institutions in securely migrating an existing role-based IAM infrastructure to an attribute-based infrastructure. Two ongoing research efforts are currently being executed, one presenting a comprehensive catalogue of key performance indicators in order to estimate the quality of existing IAM infrastructures. The other one deals with the problem of insufficient attribute quality for IAM environments based on Attribute-based Access Control (ABAC). As a result, it aims at suggesting an attribute quality framework for structured attribute quality improving initiatives. This final publication together with the four main publications mentioned above constitute the main research efforts and are intended to become part of the PhD thesis of Mr. Kunz. The thesis is in its final phase and estimated to be concluded by the end of 2017.

3.7.2 Results Achieved

Within the initiation phase of TP7, it became evident that Identity Management (IdM) covers several topics which triggered for its achieved results. Firstly, aspects on social network identity management (SIdM) were researched. One of the outcomes was the establishment of a taxonomy for social network data types (Richthammer et al. 2014). It aims at overcoming the lack of terminology for Online Social Networks regarding data types. Furthermore, we researched how end-user awareness of privacy can be increased within social networks using software (Cetto et al. 2014). The latter paper won the Best Paper Award at IDGEI 2014 and later on was awarded the prestigious Wolfgang-Heilmann-Award of 2014.

However, in contrast to SIdM, our research on recent trends within the enterprise counterpart IAM made it evident that in organizations' IAM several problems arise. Within one work, we investigated trending topics of

IAM in both research and practice, aligned them, and presented in detail which current and future challenges exist within IAM (Kunz et al. 2014). Among trends that correlate with general IT trends such as cloud-based identity management and bring your own device, the predictions on privileged user management and ABAC influenced our further research.

Based on this project, we structured our further research on IAM into five parts that later are to be integrated into a single picture, resulting in the mentioned PhD thesis within TP7. Figure 3 illustrates the five parts that TP7's research is comprising. As initial research showed, further research directions were influenced by the respective access control model (ABAC or RBAC). RBAC being the current de facto standard for access within companies, still showed shortcomings in terms of strategic maintenance. On the other hand, literature suggested that ABAC is going to become the successor of RBAC. However, a lack of research on how to establish such a change was evident. Due to experiences from existing publications on role-based IAM, our goal was to provide similar measures – adopted to attribute-based IAM. With these measures the research of TP7 aims at helping organizations firstly to be able to adopt the new access control model ABAC for their IAM while also already hinting at initial cleansing efforts and strategic maintenance undertakings.

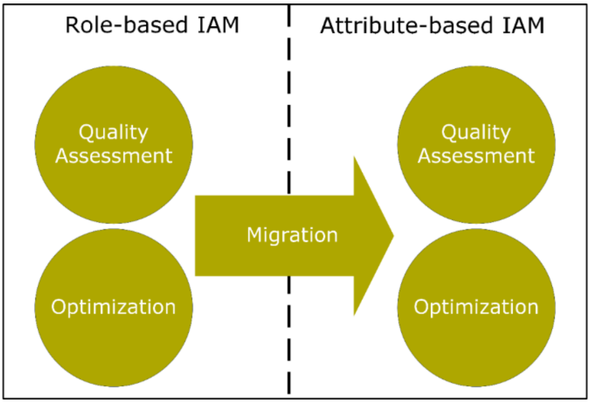

Figure 3: Main research efforts within TP7

In our first efforts, we took a closer look at Role-based IAM and currently existing research gaps. RBAC as the underlying access control model, groups employees into roles and assigns permissions to roles instead of directly granting the access rights to each employee. By doing so, the number of millions of entitlement assignments can be reduced heavily, effectively reducing complexity and improving human understandability. However, the introduction of RBAC also poses challenges if not administered properly. Outdated roles and inconstantly as well as improperly verified role assignments can lead to similar problems and role chaos.

By introducing a quality catalogue for roles and role models as a whole, we made a toolset available for responsibles that are unable to estimate their currently existing role quality (Kunz, Fuchs, Netter, et al. 2015a). Within this research, we came up with 23 criteria that role modelers can use so that they best meet the requirements for their targeted role model. Being invited to a publication of an extended version of our research, we included a dependency analysis of these quality criteria further hinting at possible side effects when increasing selected criteria (Kunz, Fuchs, Netter, et al. 2015b).

In another step, we proposed a structured process model, the Role Optimization Process Model (ROPM) (Fuchs et al. 2014) that can serve as a strategic guideline on how to revise and optimize an existing role model in complex scenarios. We proposed a four-step-procedure in which firstly respective quality goals are selected and the initial role catalogue is shrunk role by role. Afterwards, existing direct assignments of access rights to employees are to be investigated for possible role inclusion. In a last step, the role model is checked for possible hierarchical restructuring and model optimization. A practical evaluation with a real-world role model that we successfully optimized according to the client's requirements showed the practicability of our proposed procedure.

Similarly, focusing on attribute-based IAM we proposed a strategic optimization process, the Dynamic Policy Management Process (DPMP) (Hummer et al. 2015). In contrast to the ROPM, the DPMP is targeting at attribute-based IAM infrastructures. These are dominated by rules and policies thus requiring an additional focus on a more technical approach. Such environments require an initial infrastructure setup in which the systems are connected and relevant policy categories are selected. This is followed up by a data collection step in which the input data is filtered is executed. Afterwards, structured correlation and policy mining on existing access right assignments takes place in order to recommend changes to existing policies or suggest new policies. Invited to an extended version of this research, we introduced the concept of fostering KPIs within IAM. By automatically assigning risk level values to assignments of entitlements to users, we were able to identify further security risks that are candidates for inclusion within existing access control policies (Hummer et al. 2016). Again, we proved our approach within a real-world use case where we were able to identify new policies and correct existing policies for the collaborating company.

In a further effort, we came up with a migration guide that helps organizations change their predominant access control model to ABAC (Kunz, Fuchs, Hummer, et al. 2015). In this guide, we firstly aggregated building blocks of an attribute-based IAM infrastructure and the most important activities for its design. The main finding is that a two-sided approach, respecting both, attribute management and policy management activities is crucial for a successful migration. By connecting these activities in a way that presents a ubiquitous procedure, the migration guide is presented and evaluated within two real world use cases with data from a research and a healthcare institution.

In one of the last efforts within FORSEC, TP7 is currently investigating how attribute quality for attribute-based infrastructures can be improved in a structured way. We argue, that due to the highly sensitive nature and the possibly extensive consequences of attribute values within IAM, an applied framework for attribute quality management for IAM is necessary. The goal of this effort is to provide organizations with a framework that offers various methods, tools and procedures to once establish and later maintain a high-quality attribute-based IAM infrastructure. To do so, we derived critical attribute points from a generic IAM model. Building up on this model, we came up with a criteria catalogue for comparing and evaluating existing frameworks and their respective methods. This analysis paired with our expertise served as input for our attribute quality framework for attribute-based IAM. We are currently in the evaluation phase of our framework and are testing prototypical implementations with a company's data.

The last effort, currently taking place simultaneously to the establishment of an attribute quality framework is the identification and analysis of key performance indicators (KPIs) in IAM. Up to now, no structured information on how an existing IAM infrastructure, independently of its predominant access control model, is performing. Aiming at closing this gap, we firstly aggregated from both literature and experience KPIs and afterwards surveyed several IAM experts providing their expertise and views on how IAM helps their organization and which goals they are pursuing with it.

3.7.3 Contribution to FORSEC Research Alliance

Being part of the PreSTA cluster, we conducted research with both, TP10 (Economic Planning and Evaluation of IT-Security) and TP8 (Next Generation Online Trust), following the reviewers' advice on close collaboration. Through these joint works but also individually, TP7 directly influences the overall goal of FORSEC which aims at improving the security level of highly connected IT systems respecting all aspects. IAM is one of the main measures in order to help organizations to securely prepare their IT infrastructures against attacks from both, inside and outside of an organization. This can be achieved by eliminating threats stemming traditionally from manual assignment of privileges to accounts through automatic administration comprising structured approval and recertification processes. In addition, IAM poses a business enabler through higher automation and reduction of cost.

In cooperation with TP10, aiming at the economical aspect of managing privileges, we developed a decision support model. It focuses on assessing the economic value of selection of the access control model and the introduction of an Identity and Access Management System (IAMS). We fostered knowledge from the research fields of IAM and IT security investments. The outcome was a model that allows for an easy and quick but rough calculation of which of four options suits an organization that wants to evaluate an introduction of RBAC and/or an IAMS (Weishäupl, Kunz, et al. 2015).

Additionally, aiming at more securely and efficiently designing recertifications within IAM, we benefitted from the expertise of TP8 and their knowledge within the field of trust and reputation management (Richthammer et al. 2015). The new recertification model is based on the hypothesis that a constant recertification date (e.g. every 12 months) can be outperformed both in terms of security as well as the need of approval efforts by introducing trust values for employees. Integrating a temporal decay of trust level together with an outlier assessment of conspicuous assignments of privileges to employees allows for a calculation of recertification points for each employee within an organization. We demonstrated the applicability within a fictional use case based on data derived from a real-world example.

Another effort together with TP8 was the investigation of advisor attack scenarios in eCommerce reputation systems. Clustering mechanisms fostered from the field of role mining were integrated into their tool so that the visualization could be presented in a human understandable way. By using this tool, users are supported in identifying so called collusion attacks in reputation systems (Sänger, Richthammer, Kunz, et al. 2015).

3.7.4 Beyond FORSEC

TP7 provided research results mainly useful within the IAM community. By researching innovative and new procedures and tools for challenging issues within IAM, the publications of TP7 can be a starting point for further research regarding the readiness of IAM for ABAC, measures for introduction of ABAC and a holistic quality perspective on IAM. Up to now, quality is addressed only to a small amount within IAM, let alone maturity levels. This concept ensures not only a high-quality attribute infrastructure but also efficiently and effortlessly allows the management of access of identities to systems and resources. Researchers in the field of IAM and researchers within access control in particular in ABAC should work together in order to come up with more suitable solutions that have the holistic perspective on access within organizations.

Other research that can be now initiated with results from TP7 is research regarding the monitoring and lifecycle of attributes used for IAM. Companies typically undergo a large amount of attribute value changes each day. However there is only little attention to how these values change and what controls need to be established to make sure such transitions are safe. For instance, by directly connecting a company's security of authorizations to employees' attribute values, management processes for updating and outdating such values should be established. Besides, recertifications that currently are only in place for role-based IAM infrastructures need to become more fine-granular and targeted at attributes instead of roles.

Other effects of TP7 are that researchers from the DINGFEST project that had been initiated within FORSEC are now taking our results into consideration - one of the two outstanding publications has been initiated together with them, giving them deeper insight into our results so that they can make use of them within the new project.

3.7.5 Publications

Fuchs et al. (2014): Role Model Optimization for Secure Role-based Identity Management	
Abstract	In the recent past, the application of role-based access control for streamlining Identity and Access Management in organizations has gained significant importance in research and practice. After the initial setup of a role model, the central challenge is its operative management and strategic maintenance. In practice, organizations typically struggle with a high number of potentially outdated and erroneous role definitions leading to security vulnerabilities and compliance violations. Applying a process-oriented approach for assessing and optimizing role definitions is mandatory to keep a role model usable and up to date. Existing research on role system maintenance only provides a limited technical perspective without focusing on the required guidance and applicability in practice. This paper closes the existing gap by proposing ROPM, a structured Role Optimization Process Model for improving the quality of existing role definitions. Based on comprehensive tool support it automates role optimization activities and integrates both, a technical as well as a business-oriented perspective. It is based on the iterative application of role cleansing and role model extension activities in order to reduce erroneous role definitions and (re-)model roles according to organizational requirements. In order to underline applicability, this paper provides a naturalistic evaluation based on real-life data.
Citation	Fuchs, L., Kunz, M., and Pernul, G. 2014. “Role Model Optimization for Secure Role-based Identity Management,” in *Proceedings of the 22nd European Conference on Information Systems (ECIS 2014)*, Tel Aviv, Israel.
URL	http://aisel.aisnet.org/ecis2014/proceedings/track14/7/

Hummer et al. (2015): Advanced Identity and Access Policy Management using Contextual Data	
Abstract	Due to compliance and IT security requirements, company-wide Identity and Access Management within organizations has gained significant importance in research and practice over the last years. Companies aim at standardizing user management policies in order to reduce administrative overhead and strengthen IT security. Despite of its relevance, hardly any supportive means for the automated detection and refinement as well as management of policies are available. As a result, policies outdate over time, leading to security vulnerabilities and inefficiencies. Existing research mainly focuses on policy detection without providing the required guidance for policy management. This paper closes the existing gap by proposing a Dynamic Policy Management Process which structures the activities required for policy management in Identity and Access Management environments. In contrast to current approaches it fosters the consideration of contextual user management data for policy detection and refinement and offers result visualization techniques that foster human understanding. In order to underline its applicability, this paper provides a naturalistic evaluation based on real-life data from a large industrial company.
Citation	Hummer, M., Kunz, M., Fuchs, L., and Pernul, G. 2015. "Advanced Identity and Access Policy Management using Contextual Data," in *Proceedings of the 10th International Conference on Availability, Reliability and Security (ARES 2015)*, Toulouse, France.
URL	http://ieeexplore.ieee.org/document/7299897/

Hummer et al. (2016): Adaptive Identity and Access Management - Contextual Data based Policies	
Abstract	Due to compliance and IT security requirements, company-wide identity and access management within organizations has gained significant importance in research and practice over the last years. Companies aim at standardizing user management policies in order to reduce administrative overhead and strengthen IT security. These policies provide the foundation for every identity and access management system no matter if poured into IT systems or only located within responsible identity and access management (IAM) engineers' mind. Despite its relevance, hardly any supportive means for the automated detection and refinement as well as management of policies are available. As a result, policies outdate over time, leading to security vulnerabilities and inefficiencies. Existing research mainly focuses on policy detection and enforcement without providing the required guidance for policy management nor necessary instruments to enable policy adaptibility for today's dynamic IAM. This paper closes the existing gap by proposing a dynamic policy management process which structures the activities required for policy management in identity and access management environments. In contrast to current approaches, it utilizes the consideration of contextual user management data and key performance indicators for policy detection and refinement and offers result visualization techniques that foster human understanding. In order to underline its applicability, this paper provides an evaluation based on real-life data from a large industrial company.
Citation	Hummer, M., Kunz, M., Fuchs, L., and Pernul, G. 2016. “Adaptive Identity and Access Management - Contextual Data based Policies,” EURASIP Journal on Information Security (19).
URL	https://link.springer.com/article/10.1186/s13635-016-0043-2

Kunz et al. (2014): Analyzing Recent Trends in Enterprise Identity Management	
Abstract	Recent data breaches caused by highly-privileged insiders (e.g. the NSA/Snowden case) as well as the proliferation of mobile and cloud applications in enterprises imposes new challenges for Identity Management. To cope with these challenges, business analysts have predicted a variety of trends for enterprise Identity Management. In this paper, we conduct a thorough literature analysis to examine to which extent the scientific community seizes upon these trends and identify major research areas therein. Results show that despite the analysts' predictions, research stagnates for attribute-based access control and privileged user management, while for cloud-based IdM and bring your own device it corresponds to the analysts' forecast.
Citation	Kunz, M., Hummer, M., Fuchs, L., Netter, M., and Pernul, G. 2014. "Analyzing Recent Trends in Enterprise Identity Management," in *Proceedings of the 1st Workshop on Security in Highly Connected IT Systems (SHCIS 2014)*, Munich, Germany.
URL	http://ieeexplore.ieee.org/document/6974861/

Kunz, Fuchs, Hummer, et al. (2015): Introducing Dynamic Identity and Access Management in Organizations	
Abstract	Efficient and secure management of access to resources is a crucial challenge in today's corporate IT environments. During the last years, introducing company-wide Identity and Access Management (IAM) infrastructures building on the Role-based Access Control (RBAC) paradigm has become the de facto standard for granting and revoking access to resources. Due to its static nature, the management of role-based IAM structures, however, leads to increased administrative efforts and is not able to model dynamic business structures. As a result, introducing dynamic attribute-based access privilege provisioning and revocation is currently seen as the next maturity level of IAM. Nevertheless, up to now no structured process for incorporating Attribute-based Access Control (ABAC) policies into static IAM has been proposed. This paper closes the existing research gap by introducing a novel migration guide for extending static IAM systems with dynamic ABAC policies. By means of conducting structured and tool-supported attribute and policy management activities, the migration guide supports organizations to distribute privilege assignments in an application-independent and flexible manner. In order to show its feasibility, we provide a naturalistic evaluation based on two real-world industry use cases.
Citation	Kunz, M., Fuchs, L., Hummer, M., and Pernul, G. 2015. “Introducing Dynamic Identity and Access Management in Organizations,” in *Proceedings of the 11th International Conference on Information Systems Security (ICISS 2015),* Kolkata, India.
URL	https://link.springer.com/chapter/10.1007/978-3-319-26961-0_9

Kunz, Fuchs, Netter, et al. (2015b): How to discover High-quality Roles? A Survey and Dependency Analysis of Quality Criteria in Role Mining	
Abstract	Roles have evolved into the de facto standard for access control in Enterprise Identity Management. However, companies struggle to develop and maintain a role-based access control state. For the initial role deployment, role mining is widely used. Due to the high number and complexity of available role mining algorithms, companies fail to perceive which is selected best according to their needs. Furthermore, requirements on the composition of roles such as reduction of administration cost are to be taken into account in role development. In order to give them guidance, in this paper we aggregate existing role mining approaches and classify them. For consideration of individual prerequisites we extract quality criteria that should be met. Later on, we discuss interdependencies between the criteria to help role developers avoid unwanted side-effects and produce RBAC states that are tailored to their preferences.
Citation	Kunz, M., Fuchs, L., Netter, M., and Pernul, G. 2015b. "How to Discover High-quality Roles? A Survey and Dependency Analysis of Quality Criteria in Role Mining," Communications in Computer and Information Science (596).
URL	https://link.springer.com/chapter/10.1007/978-3-319-27668-7_4

Kunz, Fuchs, Netter, et al. (2015a): Analyzing Quality Criteria in Role-based Identity and Access Management	
Abstract	Roles have turned into the de facto standard for access control in enterprise identity management systems. However, as roles evolve over time, companies struggle to develop and maintain a consistent role model. Up to now, the core challenge of measuring the current quality of a role model and selecting criteria for its optimization remains unsolved. In this paper, we conduct a survey of existing role mining techniques and identify quality criteria inherently used by these approaches. This guides organizations during the selection of a role mining technique that matches their company-specific quality preferences. Moreover, our analysis aims to stimulate the research community to integrate quality metrics in future role mining approaches.
Citation	Kunz, M., Fuchs, L., Netter, M., and Pernul, G. 2015a. “Analyzing Quality Criteria in Role-based Identity and Access Management,” in *Proceedings of the 1st International Conference on Information Systems Security and Privacy (ICISSP 2015)*, Angers, France.
URL	http://ieeexplore.ieee.org/document/7509931/

3.8 TP8 – Next Generation Online Trust

3.8.1 Project Overview

TP8 addresses the issue of trust in highly connected IT systems. Due to the uncertainty of such environments, trust has become a key factor for successfully facilitating interactions and reducing the perceived risks. Although there are several ways for trust establishment, this project focusses on online reputation systems. Reputation systems provide a valuable mechanism to identify trustworthy players in an online environment. As early reputation systems could easily be manipulated by malicious actors, a huge variety of new computation models has been introduced in literature in the recent years. Although these novel models are indeed more robust, most of them have become highly complex and non-transparent to the end-user leading to a loss of understandability and a decreasing level of trust in the reputation system itself. To mitigate this weakness, TP8 introduces the concept of interactive reputation systems. Interactive reputation systems combine the dynamic configuration of computation models by the user with the transparent presentation of reputation data using interactive visualizations in the user interface. Outcomes of the experiments show that through involving the user in reputation assessment, the detection of malicious behavior can be notably enhanced whilst keeping the system's transparency.

3.8.2 Results Achieved

The generic process of a traditional reputation system includes three main phases, namely collection, computation and presentation as depicted in Figure 4. In the collection phase, the reputation system gathers all information necessary (such as ratings or textual feedback). In the second phase, the computation, one or several reputation values are computed based on a specific computation model (such as the average rating value). Finally, all outcomes are presented to the end user in the presentation phase.

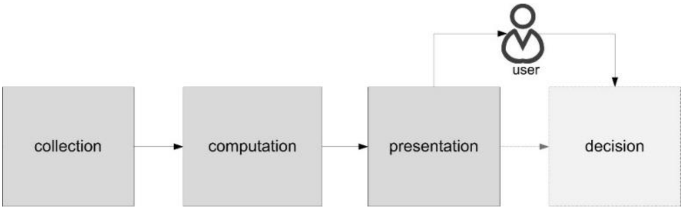

Figure 4: Generic process of a traditional reputation system

The process model of a traditional reputation systems only allows a one-way information flow from collection to presentation, respectively decision. In respect of the project goal, namely to increase the robustness of reputation systems against attacks as well as to enhance the transparency and understandability of the computation, this project explored the possibility to involve the user in reputation assessment through user interaction. This problem included three main issues: (1) How must the computation be adapted to enable interactive reputation assessment? (2) How can reputation data be transparently presented to involve the user in reputation assessment through interaction? (3) How can interactive computation and interactive presentation be integrated?

1. How must the computation be adapted to enable interactive reputation assessment?

In the first step, the computation and its components have been analyzed under the main topic "reusability" for online reputation systems. The basic idea was that reputation systems try to reflect the human reasoning process of assessing reputation and may thus be broken down to single functional building blocks each representing one factor usually influencing the decision making. These blocks could then again be individually composed to a fully functional computation engine by the user. Here, the three generic steps "filtering", "weighting" and "aggregation" could be identified based on a huge variety of different computation models of reputation systems introduced in literature. With reference to these three abstract steps, a hierarchical component taxonomy of computation components used in online reputation systems was developed. Overall – besides the three primary classes – the taxonomy includes 14 secondary component classes covering 23 component terms and 29 sub-sets. Thereby, the single components and sub-sets represent examples of how the functional blocks could be implemented. In order to encourage reusability, all components were not only named but also described on a conceptual level and stored as web services on implementation level in a component repository (http://trust.bayforsec.de/ngot/). All details on the analysis of the computation, the structure of the component repository, and the corresponding evaluation are documented in Paper (Sänger and Pernul 2014a) and Paper (Sänger, Richthammer, and Pernul 2015).

The computation components identified and implemented in the component repository in the previous step build the basis for the dynamic and flexible composition of computation engines through user interaction. In Paper (Sänger, Richthammer, Kremser, et al. 2015), we introduced a selection & composition framework that supports a user/developer in combining single components to a fully functional computation engine. The framework allows to make a selection of computation components and define all necessary setting parameters via a graphical user interface.

Thirdly, in view of the need for increased robustness, the proper usage of the component repository for attack defense was analyzed in Paper (Sänger, Richthammer, Rösch, et al. 2015). This involves a structured overview of attacks towards reputation systems, a definition of an attack taxonomy, and an allocation of reusable components stored in the component repository as example implementations of a defense technique. Details on all attacking patterns can be found online in the component repository (http://trust.bayforsec.de/ngot/).

2. How can reputation data be transparently presented to involve the user in reputation assessment through interaction?

The second part of the work in the context of TP8 addresses the transparent presentation of reputation data. In view of the overall vision – to involve the user in reputation assessment through user interaction – interactive visualizations provide a promising alternative to numerical reputation values displayed in static feedback profiles. While visualizations enable to display a much wider range of information in one view, interaction techniques allow to dive into the data. To this end, the interactive visualization of reputation data in electronic marketplaces was analyzed (this basic setup can be transferred to further environments). Based on the classification of reputation data by data types, 10 interactive visualizations have been considered. While Paper (Sänger and Pernul 2014b) and Paper (Sänger and Pernul 2016) deal with the interactive presentation of feedback and corresponding transaction context, Paper (Sänger, Richthammer, Kunz, et al. 2015) focuses on the visualization of feedback relations between different actors. Overall, three prototypes have been implemented displaying both real-world and simulated data.

Besides enabling the transparent presentation of a large amount of information in one integrated view, interactive visualizations offer a further advantage. The available interaction techniques allow to dive into the data and reveal relations and coherences. As attacks on reputation systems are characterized by a specific symptomatology, each attack shows a particular pattern. These patterns can be detected through user interaction. Paper (Sänger and Pernul 2014b), Paper (Sänger, Richthammer, Kunz, et al. 2015), and Paper (Sänger and Pernul 2016) therefore demonstrate how attacks on reputation systems can be detected using real-world and simulated data.

3. How can interactive computation and interactive presentation be integrated?

In the third step, interactive computation and interactive presentation have been integrated to the overall concept of an interactive reputation system. Reviewing the process model of a classical reputation system (Figure 4), a flow from presentation to computation necessary for user interaction is currently not considered. The classical process model was therefore extended by a bilateral flow between the computation and the presentation. As the dynamic configuration of the computation is carried out through interaction via the graphical user interface (GUI), an additional bilateral flow between the presentation (user interface) and the user was added. Figure 5 depicts the novel process model of interactive reputation systems. Based on the data collected in the first phase, both raw data (such as a list of textual reviews) and models of the data (such as reputation values or abstract visualizations) are presented to the user in the GUI. Interacting with the presentation layer, the user may dive into the data, adapt the computation models, derive insights and in this way come to a decision. Paper (Sänger and Pernul 2017) gives an overview of interactive reputation systems.

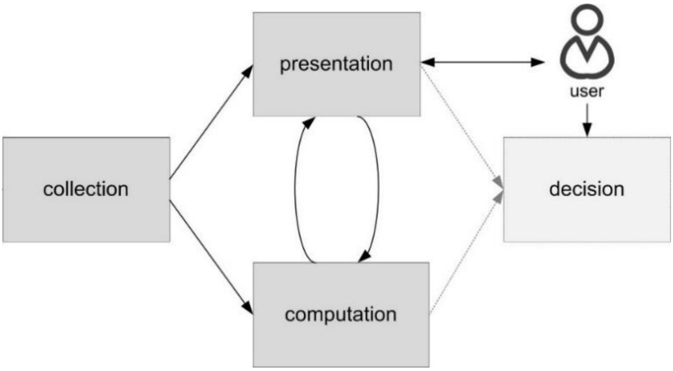

Figure 5: Generic process of an interactive reputation system

The evaluations of this novel approach to enhance the robustness whilst keeping transparency of reputation systems involved scenario analyses and case studies based on real-world and simulated data. Here, we could demonstrate that attacks can be reliably and transparently detected through user interaction (Sänger and Pernul 2014b, 2016; Sänger, Richthammer, Kunz, et al. 2015). Moreover, it remained to show that an interactive approach may indeed enhance the user's understanding of reputation data. Therefore, a user study was conducted as final evaluation that measured the user's ability to detect attacks as well as the user's understanding of reputation data. The user study was carried out in collaboration with TP6 and UCL (University College London). The study involved a controlled between-subject experiment with 40 German and 41 UK participants who had to solve four cases. In each case, participants were given the task to buy a specific item from one of two available sellers. They had to compare the two seller feedback profiles and give a preference. One of the sellers (referred to as the malicious seller) showed a discriminating behavior for one context attribute, while the other seller (referred to as the honest seller) behaved consistently with respect to the entire transaction context. The treatment group used a novel interface involving an interactive parallel coordinates visualization of reputation data (interactive reputation system) while the control group used a static eBay-like feedback profile.

Results show that participants using the novel interactive interface were significantly better able to detect malicious behavior with an overall detection accuracy of 77% versus 56% with the old interface. Only a small share of 7% decided to buy from the malicious seller as opposed to 30% in the old interface condition, leading to an increase of 178% in terms of robustness. Further results of the study revealed that through presenting all reputation data in an interactive interface, the user's understanding could be significantly enhanced (Sensemaking Score was 84% higher) despite no trade-offs in usability. All results were documented in Paper (Sänger et al. 2016).

3.8.3 Contribution to FORSEC Research Alliance

TP8 was involved both in cluster PreSTA and cluster STAR. Therefore, it served as a transfer project for the TPs of the respective clusters.

Within PreSTA, TP8 worked on several issues in joint efforts with TP7. For example, focusing on online social networks as an important example for modern highly connected systems, a user-centric categorization of data types found on these platforms was developed in collaboration with TP7 (Richthammer et al. 2014). This supported users in understanding which data types require a particular high degree of trust (in other users as well as in the service provider) when disclosing them, thus promoting awareness, privacy, and identity theft prevention. In a related project, also in collaboration with TP7, TP8 made use of a gamification approach to raise the awareness for the presentation of an individual's different identity assets on online social network platforms. Thereby, this project sought to help users establish trust in the platform as well as in the provider by giving them the possibility to playfully check their visibility settings and providing advice on this topic. The resulting application called "Friend Inspector" was awarded with the Wolfgang Heilmann Preis of the Integrata Stiftung, and the corresponding publication won a best paper award (Cetto et al. 2014). Turning to e-commerce platforms as another major example for highly connected systems, the behavior of digital identities in advisor attack scenarios was examined. In particular, clustering techniques to detect colluding adversaries were employed. Since such techniques are also used in the research field of role mining, this was another source of collaboration with TP7 (Sänger, Richthammer, Kunz, et al. 2015).

Regarding STAR, TP8 mainly contributed with research on building and analyzing ad-hoc trust relationships between unknown actors. One of the key contributions of TP8 to STAR is the provision of a method for detecting colluding adversaries on highly connected platforms. Results showed that this detection method worked well for collusion attack scenarios on e-commerce platforms such as eBay and Epinions. Moreover, TP8 integrated the "parallel coordinates visualization technique" into real eBay feedback profiles in order to assess the suitability of this technique as an active defense mechanism in the response phase. The usability and usefulness of this concept were evaluated in collaboration with TP6 (Sänger et al. 2016).

3.8.4 Beyond FORSEC

Despite our proposals being generic in theory, we have developed most of our use cases for online social networks and e-commerce platforms so far. Therefore, one part of our future research efforts focuses on using the insights gained in connection with our work in FORSEC for other application areas. In particular, we plan on examining relevant use cases of reputation systems for highly connected IT systems in the smart city environment in collaboration with the staff members of TP2 and TP9 (or their successors at the University of Passau). As a starting point of this collaboration, we envision a smart library with a reputation-based desk reservation system. While the development of the reputation system concept will be based on the work of TP8, the technical infrastructure necessary for collecting suitable input data will be implemented on top of the SERIOS platform developed in TP9.

In addition to using our trust and reputation-related insights for other application areas, we are also concerned with transferring knowledge to related research fields. We are particularly interested in recommender systems (as another kind of decision support system) because they are often applied in similar online application scenarios or even in connection with reputation systems. First, we already investigated the enhancement of recommender systems with reputation data. A systematic literature review revealed the growing interest in so-called reputation-enhanced recommender systems especially in the recent years (Richthammer et al. 2017). Second, we plan on analyzing possible classifications and visualizations of recommender systems data following the approach we used for reputation systems data (Sänger and Pernul 2016). And third, our work on interactive reputation systems clearly demonstrated the importance of considering transaction context in online environments. We want to build on the knowledge we gained in this regard to contribute to the area of context-aware recommender systems, which has received increasing attention in the recent years.

3.8.5 Publications

Cetto et al. (2014): Friend Inspector: A Serious Game to Enhance Privacy Awareness in Social Networks	
Abstract	Currently, many users of Social Network Sites are insufficiently aware of who can see their shared personal items. Nonetheless, most approaches focus on enhancing privacy in Social Networks through improved privacy settings, neglecting the fact that privacy awareness is a prerequisite for privacy control. Social Network users first need to know about privacy issues before being able to make adjustments. In this paper, we introduce Friend Inspector, a serious game that allows its users to playfully increase their privacy awareness on Facebook. Since its launch, Friend Inspector has attracted a significant number of visitors, emphasising the need for better tools to understand privacy settings on Social Networks.
Citation	Cetto, A., Netter, M., Pernul, G., Richthammer, C., Riesner, M., Roth, C., and Sänger, J. 2014. "Friend Inspector: A Serious Game to Enhance Privacy Awareness in Social Networks," in *Proceedings of the 2nd International Workshop on Intelligent Games for Empowerment and Inclusion (IDGEI 2014),* Haifa, Israel.

Hassan et al. (2014): SoDA: Dynamic Visual Analytics of Big Social Data	
Abstract	In this work we apply dynamic visual analytics on big social data by the example of microblogs from Twitter. Thereby, we address current challenges like real-time analytics as well as analyses of unstructured data. To this end, we propose SoDA - a concept enabling the integrated analysis of the dimensions: message, location and time. Furthermore, we introduce a novel design for tag cloud visualizations, the weighted tag network, offering enhanced semantic insights. All concepts are fully implemented and evaluated by a comprehensive software prototype in different application scenarios.
Citation	Hassan, S., Sänger, J., and Pernul, G. 2014. “SoDA: Dynamic Visual Analytics of Big Social Data,” in *Proceedings of the 1st International Conference on Big Data and Smart Computing (BigComp 2014),* Bangkok, Thailand.
URL	https://www.computer.org/csdl/proceedings/bigcomp/2014/3919/00/06741433-abs.html

Richthammer et al. (2014): Taxonomy of Social Network Data Types	
Abstract	Online social networks (OSNs) have become an integral part of social interaction and communication between people. Reasons include the ubiquity of OSNs that is offered through mobile devices and the possibility to bridge spatial and temporal communication boundaries. However, several researchers have raised privacy concerns due to the large amount of user data shared on OSNs. Yet, despite the large body of research addressing OSN privacy issues, little differentiation of data types on social network sites is made and a generally accepted classification and terminology for such data is missing. The lack of a terminology impedes comparability of related work and discussions among researchers, especially in the case of privacy implications of different data types. To overcome these shortcomings, this paper develops a well-founded terminology based on a thorough literature analysis and a conceptualization of typical OSN user activities. The terminology is organized hierarchically resulting in a taxonomy of data types. The paper furthermore discusses and develops a metric to assess the privacy relevance of different data types. Finally, the taxonomy is applied to the five major OSNs to evaluate its generalizability.
Citation	Richthammer, C., Netter, M., Riesner, M., Sänger, J., and Pernul, G. 2014. “Taxonomy of Social Network Data Types,” EURASIP Journal on Information Security (11).
URL	https://link.springer.com/article/10.1186/s13635-014-0011-7

Richthammer et al. (2015): Dynamic Trust-based Recertifications in Identity and Access Management	
Abstract	Security compliance has become an important topic for medium- and large-sized companies in the recent years. In order to fulfill all requirements legally imposed, high quality identity management – particularly with respect to correct and consistent access control – is essential. In this context, the concept of recertification has proven itself to maintain the quality and correctness of access rights over a long period of time. In this paper, we show how the traditional recertification concept can be notably enhanced through involving the notion of trust. We thereto propose a trust-based recertification model and demonstrate its benefits by means of a realistic use case. Our dynamic concept can help to better spread the recertification overhead compared to the traditional approach with fixed periods. Furthermore, it aids in the identification of risky employees.
Citation	Richthammer, C., Kunz, M., Sänger, J., Hummer, M., and Pernul, G. 2015. “Dynamic Trust-based Recertifications in Identity and Access Management,” in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015)*, Vienna, Austria.

Richthammer and Pernul (2016): Explorative Analysis of Recommendations Through Interactive Visualization	
Abstract	Even though today's recommender algorithms are highly sophisticated, they can hardly take into account the users' situational needs. An obvious way to address this is to initially inquire the users' momentary preferences, but the users' inability to accurately state them upfront may lead to the loss of several good alternatives. Hence, this paper suggests to generate the recommendations without such additional input data from the users and let them interactively explore the recommended items on their own. To support this explorative analysis, a novel visualization tool based on treemaps is developed. The analysis of the prototype demonstrates that the interactive treemap visualization facilitates the users' comprehension of the big picture of available alternatives and the reasoning behind the recommendations. This helps the users get clear about their situational needs, inspect the most relevant recommendations in detail, and finally arrive at informed decisions.
Citation	Richthammer, C., and Pernul, G. 2016. "Explorative Analysis of Recommendations Through Interactive Visualization," in *Proceedings of the 17th International Conference on Electronic Commerce and Web Technologies (EC-Web 2016),* Porto, Portugal.
URL	https://link.springer.com/chapter/10.1007/978-3-319-53676-7_4

Richthammer et al. (2017): Interactive Visualization of Recommender Systems Data	
Abstract	Recommender systems provide a valuable mechanism to address the information overload problem by reducing a data set to the items that may be interesting for a particular user. While the quality of recommendations has notably improved in the recent years, the complex algorithms in use lead to high non-transparency for the end user. We propose the usage of interactive visualizations for presenting recommendations. By involving the user in the information reduction process, the quality of recommendations could be enhanced whilst keeping the system's transparency. This work gives first insights by analyzing recommender systems data and matching them to suitable visualization and interaction techniques. The findings are illustrated by means of an example scenario based on a typical real-world setting.
Citation	Richthammer, C.; Sänger, J.; Pernul, G. 2017. Interactive Visualization of Recommender Systems Data. In: *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems (SHCIS 2017).* Neuchâtel, Switzerland.

Richthammer et al. (2017): Reputation-Enhanced Recommender Systems	
Abstract	Recommender systems are pivotal components of modern Internet platforms and constitute a well-established research field. By now, research has resulted in highly sophisticated recommender algorithms whose further optimization often yields only marginal improvements. This paper goes beyond the commonly dominating focus on optimizing algorithms and instead follows the idea of enhancing recommender systems with reputation data. Since the concept of reputation-enhanced recommender systems has attracted considerable attention in recent years, the main aim of the paper is to provide a comprehensive survey of the approaches proposed so far. To this end, existing work are identified by means of a systematic literature review and classified according to carefully considered dimensions. In addition, the resulting structured analysis of the state of the art serves as a basis for the deduction of future research directions.
Citation	Richthammer, C., Weber, M., and Pernul, G. 2017. "Reputation-Enhanced Recommender Systems," in *Proceedings of the 11th IFIP WG 11.11 International Conference on Trust Management (IFIPTM 2017),* Gothenburg, Sweden.
URL	https://link.springer.com/chapter/10.1007/978-3-319-59171-1_13

Sänger and Pernul (2014a): Reusability for Trust and Reputation Systems	
Abstract	Reputation systems have been extensively explored in various disciplines and application areas. A problem in this context is that the computation engines applied by most reputation systems available are designed from scratch and rarely consider well established concepts and achievements made by others. Thus, approved models and promising approaches may get lost in the shuffle. In this work, we aim to foster reuse in respect of trust and reputation systems by providing a hierarchical component taxonomy of computation engines which serves as a natural framework for the design of new reputation systems. In order to assist the design process we, furthermore, provide a component repository that contains design knowledge on both a conceptual and an implementation level.
Citation	Sänger, J., and Pernul, G. 2014a. “Reusability for Trust and Reputation Systems,” in *Proceedings of the 8th IFIP WG 11.11 International Conference on Trust Management (IFIPTM 2014),* Singapore.
URL	https://link.springer.com/chapter/10.1007/978-3-662-43813-8_3

Sänger and Pernul (2014b): Visualizing Transaction Context in Trust and Reputation Systems	
Abstract	Transaction context is an important aspect that should be taken into account for reputation-based trust assessment, because referrals are bound to the situation-specific context in which they were created. The non-consideration of transaction context may cause several threats such as the value imbalance problem. Exploiting this weakness, a seller can build high reputation by selling cheap products while cheating on the expensive ones. In the recent years, multiple approaches have been introduced that address this challenge. All of them chose metrics leading to numerical reputation values. These values, however, are non-transparent and quite hard to understand for the end-user. In this work, in contrast, we combine reputation assessment and visual analytics to provide an interactive visualization of multivariate reputation data. We thereby allow the user to analyze the data sets and draw conclusions by himself. In this way, we enhance transparency, involve the user in the evaluation process and as a consequence increase the users' trust in the reputation system.
Citation	Sänger, J., and Pernul, G. 2014b. “Visualizing Transaction Context in Trust and Reputation Systems,” in *Proceedings of the 9th International Conference on Availability, Reliability and Security (ARES 2014),* Fribourg, Switzerland.
URL	http://ieeexplore.ieee.org/document/6980268/

Sänger et al. (2014): Trust and Big Data: A Roadmap for Research	
Abstract	We are currently living in the age of Big Data coming along with the challenge to grasp the golden opportunities at hand. This mixed blessing also dominates the relation between Big Data and trust. On the one side, large amounts of trust-related data can be utilized to establish innovative data-driven approaches for reputation-based trust management. On the other side, this is intrinsically tied to the trust we can put in the origins and quality of the underlying data. In this paper, we address both sides of trust and Big Data by structuring the problem domain and presenting current research directions and inter-dependencies. Based on this, we define focal issues which serve as future research directions for the track to our vision of Next Generation Online Trust within the FORSEC project.
Citation	Sänger, J., Richthammer, C., Hassan, S., and Pernul, G. 2014. "Trust and Big Data: A Roadmap for Research," in *Proceedings of the 1st Workshop on Security in Highly Connected IT Systems (SHCIS 2014)*, Munich, Germany.
URL	http://ieeexplore.ieee.org/document/6974862/

Sänger, Richthammer, Kremser, et al. (2015): Personalized Composition of Trustful Reputation Systems	
Abstract	The vast amount of computation techniques for reputation systems proposed in the past has resulted in a need for a global online trust repository with reusable components. In order to increase the practical usability of such a repository, we propose a software framework that supports the user in selecting appropriate components and automatically combines them to a fully functional computation engine. On the one hand, this lets developers experiment with different concepts and move away from one single static computation engine. On the other hand, our software framework also enables an explorative trust evaluation through user interaction. In this way, we notably increase the transparency of reputation systems. To demonstrate the practical applicability of our proposal, we present realistic use cases and describe how it would be employed in these scenarios.
Citation	Sänger, J., Richthammer, C., Kremser, A., and Pernul, G. 2015. "Personalized Composition of Trustful Reputation Systems," in *Proceedings of the 29th Annual IFIP WG 11.3 Working Conference on Data and Applications Security and Privacy (DBSec 2015),* Fairfax, VA.
URL	https://link.springer.com/chapter/10.1007/978-3-319-20810-7_13

Sänger, Richthammer, Kunz, et al. (2015): Visualizing Unfair Ratings in Online Reputation Systems	
Abstract	Reputation systems provide a valuable method to measure the trustworthiness of sellers or the quality of products in an e-commerce environment. Due to their economic importance, reputation systems are subject to many attacks. A common problem are unfair ratings which are used to unfairly increase or decrease the reputation of an entity. Although being of high practical relevance, unfair rating attacks have only rarely been considered in literature. The few approaches that have been proposed are furthermore quite non-transparent to the user. In this work, we employ visual analytics to identify colluding digital identities. The ultimate benefit of our approach is the transparent revelation of the true reputation of an entity by interactively using both endogenous and exogenous discounting methods. We thereto introduce a generic conceptual design of a visual analytics component that is independent of the underlying reputation system. We then describe how this concept was implemented in a software prototype. Subsequently, we demonstrate its proper functioning by means of an empirical study based on two real-world datasets from eBay and Epinions. Overall, we show that our approach notably enhances transparency, bares an enormous potential and might thus lead to substantially more robust reputation systems and enhanced user experience.
Citation	Sänger, J., Richthammer, C., Kunz, M., Meier, S., and Pernul, G. 2015. “Visualizing Unfair Ratings in Online Reputation Systems,” in *Proceedings of the 23rd European Conference on Information Systems (ECIS 2015),* Münster, Germany.
URL	http://aisel.aisnet.org/ecis2015_cr/159/

Sänger, Richthammer, and Pernul (2015): Reusable Components for Online Reputation Systems	
Abstract	Reputation systems have been extensively explored in various disciplines and application areas. A problem in this context is that the computation engines applied by most reputation systems available are designed from scratch and rarely consider well established concepts and achievements made by others. Thus, approved models and promising approaches may get lost in the shuffle. In this work, we aim to foster reuse in respect of trust and reputation systems by providing a hierarchical component taxonomy of computation engines which serves as a natural framework for the design of new reputation systems. In order to assist the design process we, furthermore, provide a component repository that contains design knowledge on both a conceptual and an implementation level. To evaluate our approach we conduct a descriptive scenario-based analysis which shows that it has an obvious utility from a practical point of view. Matching the identified components and the properties of trust introduced in literature, we finally show which properties of trust are widely covered by common models and which aspects have only rarely been considered so far.
Citation	Sänger, J., Richthammer, C., and Pernul, G. 2015. “Reusable Components for Online Reputation Systems,” Journal of Trust Management (2:5).
URL	https://link.springer.com/article/10.1186/s40493-015-0015-3

Sänger, Richthammer, Rösch, et al. (2015): Reusable Defense Components for Online Reputation Systems	
Abstract	Attacks on trust and reputation systems (TRS) as well as defense strategies against certain attacks are the subject of many research papers. Although proposing valuable ideas, they all exhibit at least one of the following major shortcomings. Firstly, many researchers design defense mechanisms from scratch and without reusing approved ideas. Secondly, most proposals are limited to naming and theoretically describing the defense mechanisms. Another issue is the inconsistent denomination of attacks with similar characteristics among different researchers. To address these shortcomings, we propose a novel taxonomy of attacks on TRS focusing on their general characteristics and symptomatology. We use this taxonomy to assign reusable, clearly described and practically implemented components to different classes of attacks. With this work, we aim to provide a basis for TRS designers to experiment with numerous defense mechanisms and to build more robust systems in the end.
Citation	Sänger, J., Richthammer, C., Rösch, A., and Pernul, G. 2015. “Reusable Defense Components for Online Reputation Systems,” in *Proceedings of the 9th IFIP WG 11.11 International Conference on Trust Management (IFIPTM 2015)*, Hamburg, Germany.
URL	https://link.springer.com/chapter/10.1007/978-3-319-18491-3_15

Sänger et al. (2016): Look Before You Leap: Improving the Users' Ability to Detect Fraud in Electronic Marketplaces	
Abstract	Reputation systems in current electronic marketplaces can easily be manipulated by malicious sellers in order to appear more reputable than appropriate. We conducted a controlled experiment with 40 UK and 41 German participants on their ability to detect malicious behavior by means of an eBay-like feedback profile versus a novel interface involving an interactive visualization of reputation data. The results show that participants using the new interface could better detect and understand malicious behavior in three out of four attacks (the overall detection accuracy 77% in the new vs. 56% in the old interface). Moreover, with the new interface, only 7% of the users decided to buy from the malicious seller (the options being to buy from one of the available sellers or to abstain from buying), as opposed to 30% in the old interface condition.
Citation	Sänger, J., Hänsch, N., Glass, B., Benenson, Z., Landwirth, R., and Sasse, M. A. 2016. “Look Before You Leap: Improving the Users' Ability to Detect Fraud in Electronic Marketplaces,” in *Proceedings of the 34th ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 2016).*
URL	https://dl.acm.org/citation.cfm?id=2858555

Sänger and Pernul (2016): TRIVIA: Visualizing Reputation Profiles to Detect Malicious Sellers in Electronic Marketplaces	
Abstract	Reputation systems are an essential part of electronic marketplaces that provide a valuable method to identify honest sellers and punish malicious actors. Due to the continuous improvement of the computation models applied, advanced reputation systems have become non-transparent and incomprehensible to the end-user. As a consequence, users become skeptical and lose their trust toward the reputation system. In this work, we are taking a step to increase the transparency of reputation systems by means of providing interactive visual representations of seller reputation profiles. We thereto propose TRIVIA - a visual analytics tool to evaluate seller reputation. Besides enhancing transparency, our results show that through incorporating the visual-cognitive capabilities of a human analyst and the computing power of a machine in TRIVIA, malicious sellers can be reliably identified. In this way we provide a new perspective on how the problem of robustness could be addressed.
Citation	Sänger, J., and Pernul, G. 2016. "TRIVIA: Visualizing Reputation Profiles to Detect Malicious Sellers in Electronic Marketplaces," Journal of Trust Management (3:5).
URL	https://link.springer.com/article/10.1186/s40493-016-0026-8

3.9 TP9 – Web Security

3.9.1 Project Overview

TP9 considers web security in the context of ubiquitous computing and the Internet of Things. For this purpose, TP9 analyses how evolving web technologies can be used to support smart home and smart city environments in a secure way, and how these technologies need to be changed and adapted. In the first phase of the project, we developed mechanisms for secure machine-to-machine authentication and secure session migration. This effort was based on a survey of protocols suitable for smart home environments. In a second step, we developed techniques which prevent CSRF based on enhanced security policies. We complemented these security mechanisms by enforcing control flow integrity of web applications.

The second phase of the project assumed a secure browser environment and focused on the design and implementation of an integrated and concise data-centric security model. TP9 targeted a unified security architecture which allows the consistent enforcement of fine-grained and data-centric security policies generated by or targeted for the platform services developed in TP2 and processed by the web-based services designed to efficiently interact and exploit the underlying IoT.

3.9.2 Results Achieved

The activities and results in this TP are two-pronged: On the one hand, we focused on existing security threats for web-based IoT applications executed in a browser environment and on the other hand we look at the security requirements of IoT applications which are based on web-technologies.

In the browser domain, we have developed mechanisms for secure machine-to-machine (M2M) authentication. We also designed and implemented browser support for enhanced security policies against CSRF which are based on the user's authentication status. Finally, we also developed an approach to enforce valid control flows in web applications to prevent the execution of potentially malicious execution paths which may yield confidential information or perform unauthorized transactions. In addition, we developed a concept for migrating sessions between different devices or browsers which is essential to maintain the security state established by the technologies introduced above.

In the second branch of our activities we focused on web-technologies which would enhance the data-centric approach for security in IoT-applications. This effort complements the activities in TP2 and investigates the full power of such an approach for the IoT.

For this purpose, we conducted a literature review to explore different technologies. We derived a general architecture for IoT applications and derived classes of security perimeters which can induce the application of different security mechanisms available today.

As this result mainly covers technologies which ensure confidentiality and integrity properties for data in transit from IoT devices to typical back-ends of web applications, it was fundamental for the next tasks. They will mainly focus on security mechanisms at the application layer and will assume the use of the technologies we surveyed and of technologies developed in TP2.

Currently, many commercial as well as open IoT platforms rely on web and cloud technologies. This simplifies and speeds-up application development and increases platform interoperability but decreases the control over private data. While a paradigm shift to data-centric security would greatly mitigate this problem, integrating data-centric security policies, their enforcement, and monitoring into typical IoT applications becomes a non-trivial task.

With this observation in mind, this TP designed a lightweight IoT platform which is based on approaches investigated in the FP7 project COMPOSE. It targets smaller gateways such as ODroid which have been used to design a SmartHome environment in TP2 or can be used in the cloud to scale to larger applications, such as SmartCities. The prototype, called SEDARI, is open source and shows the feasibility of data-centric security policies by integrating the UPFROnt policy framework which was also developed in this TP. It is an extension of the Usage Locks Policy framework developed in COMPOSE and offers flexible definition of locks, i.e. developers can define policy frameworks according to their needs, and includes several declassification strategies. Apart from an extended ULock policy language and classical policy framework components, UPFROnt defined and implements architectural enhancement, such as PAP and PIP caches, which can be integrated into IoT applications to ease the negative performance impact of our data-centric approach. UPFROnt also integrates with SERIOS, our data back end, that has been designed to manage the virtual twins of and to provide the working memory for IoT devices. SERIOS can store IoT devices and data and associates them transparently with fine-grained security policies. TP2 uses SERIOS to manage its devices, to store the data they generate, and to define appropriate policies on data and entities.

Of course, UPFROnt requires an appropriate identity management system. For this purpose, we collaborated with the H2020 project AGILE to enhance SEDARI with the AGILE IDM system. It does not only turn SEDARI into an Oauth2 server but it can also take the role of an OAuth2 provider. We consider both capabilities are essential for IoT applications.

For the development and execution of IoT application which respects data-centric security policies, we used and enhanced the popular flow editor Node-RED. First, we continued and refined our work started in COMPOSE. We integrated a user interface which supports the extended version of Usage Locks into the latest version of the Node-RED editor. It allows users to define policies for Node-RED nodes as well as devices registered with SERIOS. Policies are managed in UPFROnt. This information is used by the static workflow checker we integrated in Node-RED as well. It identifies potential security problems which may occur during runtime and informs users before deployment.

During runtime, we use the same principles as in COMPOSE to propagate security policies but adopted an asynchronous security policy processing in the new version of Node-RED. This also holds for function nodes which have been executed in JSFlow, a JavaScript interpreter guaranteeing non-interference properties. We extended JSFlow to support the enhanced version of Usage Locks developed in this TP. Additionally we investigated new methods to improve the internal policy management mechanisms of JSFlow. They induced large overheads in data structures which are common in IoT applications interacting with RESTful interfaces. An approach similar to compression has been designed and implemented in the v8 context to reduce the amount of redundant information stored in a JSFlow wrapper object. Further, a major extension to JSFlow now also allows the loading of modules, a crucial feature for code designed for the IoT and node.js. Main research challenge was the fact that module loading also allows the inclusion of native code. Thus, without policy tracking in native code, applications could simply bypass policy enforcement. For this reason we adopted a similar approach as in TP2 and use Valgrind to monitor security policy propagation in a separate process. However, instead of only tracking simple taint values, we implemented a completely new tool, PolicyGrind, which now allows policy tracking on various precision levels, i.e. with bit-, byte-, word- and variable precision. Our JSFlow-PolicyGrind combination is called POLITE.

Policy configurations induce a large number of variables in particular when IoT-devices and applications defined over them are shared. Using a workflow editor in this setting may appear to be easy but it also hides a lot of details. As a consequence, workflow configurations may contain a lot of security conflicts. As our policy framework is based on flow control policies, it improves the control of users over their data or the data produced by devices they own. However, this also induces the number of conflicts identified by our flow checker and makes fixing of such conflicts even more difficult. Considering that Node-RED has been designed for non security experts, we enhanced the workflow editor with a constraint solver which proposes reconfigurations of workflows by injecting so called security services. They perform declassification operations on data or establish additional security associations with communication end-points. In fact, this result shows the real power of our usage lock policy language as it can also guide reconfiguration or rewriting processes. We already see applications in other domains, e.g. in the Android application ecosystem or in micro-service combinations.

Finally, we are currently integrating all these features in one unified Node-RED prototype, called NEROS. The latter will be included in SEDARI and we will make it available as an open source implementation until the end of the project. We also think about offering our tools to a set of developers by instantiating it on local hardware. Currently, we only offer restricted access from inside the University at https://sedari.sec.uni-passau.de/.

For experimental purposes and case studies we also built three applications: the SmartHome, the SmartDoorSign, and the SmartLibrary.

The first application implements a smart home. It was developed in TP2 and manages its IoT devices and data in SERIOS. With this application, we currently investigate how and which security mechanisms, integrated in the SmartHome Android App, can be generalized and mapped into Node-RED services or even SEDARI security services. This would allow the reuse or automated integration of new security components in similar application scenarios. The SmartHome also shows how tedios and error prone access control implementations on security sensitive data can be replaced by UPFROnt. In this way, we show a unified and integrated way to track policy-annotated information from the device level through the gateway (TP2) to sets of web-applications (TP9).

In the second demo application, we show case similar properties but enhance it with private contextual information, e.g. user location, which is combined with other private information from regular sources, such as calendars. Our SmartDoorSign demo targets the efficient use of office space at Universities and tries to increase availability and communication between university members and students. From BLE beacons we infer the location of staff members. By using Usage Lock policies, they can define when their location or other information, e.g. calendar entries, should be displayed to whom and decide on the level of detail or precision of the visible information (through declassification). This contextual information can be combined and appropriately filtered by SEDARI and be shown on displays (currently a tables) mounted at the office of a person or on an App.

The third demo was started as a collaboration with the ACTLab at the University of Passau during COMPOSE. Goal of this show case is to increase the availability of library seats. For this purpose, 28 tables of the library have been equipped with privacy preserving temperature sensors mounted on the lamps attached to each table. They are connected to a local gateway. During FORSEC the collaboration with ACTLab continued and we now accumulate library data in the policy enhanced storage of SERIOS. The Node-RED services originally defined in COMPOSE have been extended and adopted to the new Usage Lock policies and run in SEDARI. We can not only show their verification and reconfiguration but we are also able to demonstrate the impact of policy changes on the information delivered to a Cordova App. The latter has been implemented in a collaborative effort and can be used by students or library staff. Depending on the security policies and the users the application interacts with, it shows real-time occupancy data, private user information, and possibly reservation information. While being an extremely helpful playground for the development and refinement of our policy framework, this demo successfully motivates the urgent need for the security mechanisms implemented in this TP.

3.9.3 Contribution to FORSEC Research Alliance

In collaboration with TP2, we show how to shape a unified security architecture which allows the use of fine-granular, data-centric security policies in the IoT. Data items annotated with security policies can propagate from sensors, through a device executing applications on this data, to a gateway storing or forwarding the data to a web-service which associates data with a virtual twin and its policies. By interacting with this web-service, the data can then be processed by web-applications. In all these steps TP2 and TP9 were able to show how sensors, devices, gateways and web-applications interact securely and – with certain limitations and assumptions required for the realization of our prototypes – show how the user is able to define security constraints on how the data can be processed.

As such, TP9 showed how to mitigate the threats for IoT applications which are induced by the web-technologies currently used to quickly promote these applications. We complemented this effort to also show the benefits of these young technologies and investigated how they can be combined with existing theoretical frameworks to increase user governance on data and to also provide non-security experts during the design phase of IoT applications. For this purpose, we provide mechanisms which help to identify potentially misconfigured or malicious services. In fact, we go beyond detection and support the user in fixing potentially harmful service configurations or implementations. This is not only an essential defensive measure but it also simplifies security debugging and ensures service functionality. On top, we designed reference monitors which can also identify security conflicts during execution. Thus, we also facilitate the analysis of runtime security problems.

Further, with the integration of a modern identity management and authorization system we enable cross-domain interactions of client applications and users. As the principals of our data-centric policy framework are defined by this system, we provide an important basis to thoroughly and securely integrate several systems with different security requirements, a crucial property for highly interconnected systems, particularly for the IoT. Of course, we are aware that our solution requires several trust assumptions and is currently subject to several limitations. However, at the same time, together with TP2, we show the feasibility of a new way to ensure secure interaction: a data-centric security approach.

In summary, TP9 addresses two important aspects of defensive measures with the unified IoT architecture and development platform SEDARI. It offers a tool for a unified IT-security process that already prepares against potential threats to privacy critical data, increases the awareness of users and provides feasible mechanisms supporting the analysis of incidents and declassification mechanisms to guarantee policy compliance.

3.9.4 Beyond FORSEC

Until the end of FORSEC we will focus on submitting the obtained results to appropriate conferences and workshops.

After the end of the project, we plan to partially **reuse SEDARI** in the implementation of two European proposals we are currently involved in. In particular, we will investigate the trade-off between data-centric mechanisms which increase performance and reduce overhead and the security implications for these optimizations, i.e. the precision and completeness.

The **installations and demos** will help us to further investigate the feasibility of our approach, and understand user acceptance of data-centric security policies. In particular, for the SmartLibrary use case we will study how **online trust** (TP8) can be integrated in the security policy framework, and based on experiments with users, understand the impact of enforcement, based on online-trust, on the overall occupancy of the library.

The **PolicyGrind** tool developed in this TP is based on Valgrind and tracks policy information in native code. In collaboration with TP2, we have also designed a Valgrind tool for taint tracking in native code included in Android applications. In a future step, we will replace this tool with PolicyGrind in order to enhance the power of FlowCoaster. Additionally, we will focus on the security weaknesses of Valgrind and introduce additional features to prevent the attacks we have already identified during our work on PolicyGrind.

Recently and inspired by the collaboration between the ACTLab and the IT-Security Chair at the University of Passau, we discovered that a new **preference-based automation** approach for the IoT can be extended with the security constraints defined by the usage lock framework and can in fact be integrated in the Node-RED framework of SEDARI. Based on templates we will simplify and automate service composition and restrain this automation with data-centric security policies. On top, these secure compositions will be of dynamic nature, trying to satisfy user preferences as good as possible.

3.9.5 Publications

Parra et al. (2016): Addressing Data-Centric Security Requirements for IoT-Based Systems	
Abstract	Allowing users to control access to their data is paramount for the success of the Internet of Things, therefore, it is imperative to ensure it, even when data has left the users' control, e.g. shared with cloud infrastructure. Consequently, we propose several state of the art mechanisms from the security and privacy research fields to cope with this requirement. To illustrate how each mechanism can be applied, we derive a data-centric architecture providing access control and privacy guaranties for the users of IoT-based applications. Moreover, we discuss the limitations and challenges related to applying the selected mechanisms to ensure access control remotely. Also, we validate our architecture by showing how it empowers users to control access to their health data in a quantified self use case.
Citation	Parra, J. D., Schreckling, D., and Posegga, J. 2016. “Addressing Data-Centric Security Requirements for IoT-Based Systems,” in P*roceedings of the International Workshop on Secure Internet of Things (SIOT 2016),* Oslo, Norway.
URL	http://ieeexplore.ieee.org/document/7913560/

Braun, Köstler, et al. (2014): A Trusted UI for the Mobile Web	
Abstract	Modern mobile devices come with first class web browsers that rival their desktop counterparts in power and popularity. However, recent publications point out that mobile browsers are particularly susceptible to attacks on web authentication, such as phishing or clickjacking. We analyze those attacks and find that existing countermeasures from desktop computers can not be easily transfered to the mobile world. The attacks' root cause is a missing trusted UI for security critical requests. Based on this result, we provide our approach, the MobileAuthenticator, that establishes a trusted path to the web application and reliably prohibits the described attacks. With this approach, the user only needs one tool to protect any number of mobile web application accounts. Based on the implementation as an app for iOS and Android respectively, we evaluate the approach and show that the underlying interaction scheme easily integrates into legacy web applications.
Citation	Braun, B., Köstler, J., Posegga, J., and Johns, M. 2014. "A Trusted UI for the Mobile Web," in *Proceedings of the 29th IFIP International Information Security and Privacy Conference (IFIP SEC 2014).*
URL	https://link.springer.com/chapter/10.1007/978-3-642-55415-5_11

Braun, Gries, et al. (2014): Ghostrail: Ad Hoc Control-Flow Integrity for Web Applications	
Abstract	Modern web applications frequently implement complex control flows, which require the users to perform actions in a given order. Users interact with a web application by sending HTTP requests with parameters and in response receive web pages with hyperlinks that indicate the expected next actions. If a web application takes for granted that the user sends only those expected requests and parameters, malicious users can exploit this assumption by crafting harming requests. We analyze recent attacks on web applications with respect to user-defined requests and identify their root cause in the missing enforcement of allowed next user requests. Based on this result, we provide our approach, named Ghostrail, a control-flow monitor that is applicable to legacy as well as newly developed web applications. It observes incoming requests and lets only those pass that were provided as next steps in the last web page. Ghostrail protects the web application against race condition exploits, the manipulation of HTTP parameters, unsolicited request sequences, and forceful browsing. We evaluate the approach and show that it neither needs a training phase nor a manual policy definition while it is suitable for a broad range of web technologies.
Citation	Braun, B., Gries, C., Petschkuhn, B., and Posegga, J. 2014. “Ghostrail: Ad Hoc Control-Flow Integrity for Web Applications,” in *Proceedings of the 29th IFIP International Information Security and Privacy Conference (IFIP SEC 2014).*
URL	https://link.springer.com/chapter/10.1007/978-3-642-55415-5_22

Braun et al. (2015): LogSec: Adaptive Protection for the Wild Wild Web	
Abstract	Today, a Web browser is a user's gateway to a multitude of Web applications, each with its own balance between confidentiality and integrity versus cross-application content sharing. Modern Web browsers apply the same permissive security policy to all content regardless of its demand for security – a behavior that enables attacks such as cross-site request forgery (CSRF) or sidejacking. To defend against such attacks, existing countermeasures enforce overly strict policies, which expose incompatibilities with real-world Web applications. As a consequence, users get annoyed by malfunctions. In this paper, we show how browser behavior can be adapted based on the user's authentication status. The browser can enforce enhanced security policies, if necessary, and permit modern communication features, if possible. Our approach mitigates CSRF, session hijacking, sidejacking, and session fixation attacks. We present the implementation as a browser extension, named LogSec, that passively detects the user's authentication status without server-side support and is transparent for the user.
Citation	Braun, B., Pauli, K., Posegga, J., and Johns, M. 2015. "LogSec: Adaptive Protection for the Wild Wild Web," in 2015 ACM Symposium on Applied Computing (SAC 2015).
URL	https://dl.acm.org/citation.cfm?id=2695709

3.10 TP10 – Economic Planning and Evaluation of IT Security

3.10.1 Project Overview

As the use of information technology (IT) steadily increases, the economic perspective on information security gains relevance: Security incidents can lead to disruption of production and processes or data theft, which, in turn, result in economic damage, including losses in productivity and income, strategic disadvantages and loss of reputation. To avoid these damages, firms invest into various security measures that protect systems, data and processes against technical failure, damage or attacks. However, when it comes to information security investments, organizations face two key challenging tasks: making decisions with regard to priorities and budgets of investments in security countermeasures by estimating the costs and benefits of possible investments (ex ante perspective) and evaluating the effectiveness and efficiency of past investments in security countermeasures to improve future investment decisions (ex post perspective and learning).

The goal of this research project is the adoption of an economic perspective on IT security in highly-distributed environments in order (1) to develop a theoretical foundation of information security investments, (2) to suggest practically feasible decision support models and evaluation metrics, and (3) to use the developed theoretical artefacts in empirical investigations of IT security investments.

3.10.2 Results Achieved

In the following, we describe the results achieved along the three subgoals defined in the project overview. As recommended by the reviewers in their evaluation report, we put emphasis to empirical and qualitative approaches rather than quantitative concepts.

Addressing the first goal, we developed a new theoretical model emerging from a multi-theoretical perspective adopting two established theories in information systems research, the resource-based view (RBV) and the organizational learning theory (OLT). The joint application of these theories allows to (i) consider firm-, industry- and national-level effects of information security investments on organizational performance, (ii) disaggregate these effects and identify the roles of security resources and security processes in the overall IT business value generation process, and (iii) to account for changes of both IT and business environments by theorizing on learning effects in terms of adapted security strategies and security actions. We used the suggested theoretical model to develop a first comprehensive map of research (results) on IT security investments by reviewing the literature and structuring the insights along the model. We also identified previously neglected areas of IT security investments, pointing to research gaps and avenues for further research. The aforementioned contributions (Reinfelder and Weishäupl 2016; Schryen and Weishäupl 2015; Weishäupl, Yasasin, et al. 2015a, 2015b; Yasasin et al. 2017) include the following results and implications: (1) There is a substantial body of literature on the importance of considering environmental factors when investing in information security resources. However, it is yet not understood how these factors interact and jointly affect investment decisions. (2) A positive impact of both technological and human information security investments on the organizational performance of a firm has been identified and measured using different metrics. We expect that this impact is mediated through its influence on security processes and business processes as there is consensus in the literature that the causal relationship between investments in IT assets in general and the organizational performance shows such mediation effects. It is yet unknown whether and how investments in proactive and reactive information security assets differ in their effects on organizational performance when this influence is mediated through security and business processes. (3) While the existence of relationships between security process (performance) and business process (performance) and the impact of these relationships on the organizational performance is acknowledged in the literature, there is only a vague understanding of the nature of these relationships. (4) Organizations need to constantly learn from the impact that past information security investments have had on the organizational performance and adapt their long-term strategies and medium-term actions.

With regard to the second goal, we developed (1) several decision support models and (2) evaluation metrics.

(1) We developed a support model for IT security incident management to effectively assign and schedule security incidents to the members of the IT staff (Rauchecker et al. 2014). We use methods of operations research (OR) to propose an optimization model to optimally assign and schedule these incidents. By doing this, we bridged the gap between the quantitative methods of OR and the field of IT security management. Moreover, we proposed efficient solution methods: we showed the practical applicability of our approach by developing efficient solution heuristics. Numerical simulations proved that our approach improves current best practice behavior significantly.

We proposed another decision support model focusing on IT security investments in highly distributed systems based on fuzzy set theory (Yasasin et al. 2014). We used fuzzy set theory, which is an established uncertainty theory because, in practice, decision makers often face non-probabilistic uncertainty regarding budget constraints, costs and security levels within highly distributed systems. We further developed and tested a Monte Carlo heuristic to solve the optimization problem.

In addition, we developed a multi-objective optimization model to support information security investment decision making considering two conflicting objectives, the minimization of the costs of countermeasures and the maximization of the resulting security level (Weishäupl 2017). These goals are modelled as objective functions which are optimized with respect to the implementation of security controls in the presence of hard constraints on the variables. For the decision support model, the classic components of risk analysis, including assets, controls, vulnerabilities and threats, and their interdependencies were considered.

(2) For developing security evaluation metrics, we drew on argumentation theory to derive requirements that IT security metrics should fulfill (Yasasin and Schryen 2015). As a contribution, five key requirements were derived, their implications discussed and exemplarily applied to two practically used IT security metrics. IT security metrics should be (a) bounded, (b) metrically scaled, (c) reliable, valid and objective, (d) context-specific and (e) computed automatically in order to be practicable applicable. By applying the requirements to two practically used IT security metrics, we were able to demonstrate that the proposed requirements ensure objective, valid, reliable and clearly interpretable security metrics.

For the third goal, we conducted an exploratory case study with firms. We strove to understand how information security investment decisions are made and evaluated in firms and how they organizationally learn from past experience. Thus, we used the developed theoretical model to structure our interviews with the firms. These interviewed firms included consulting organizations in order to find out how security investment decisions, evaluations and learning strategies are conducted by their clients. As consulting firms often have only limited insights in the security investment management of their clients, we also conducted interviews with non-consulting firms.

A scenario-specific approach – as initially planned – was not possible because we could not recruit a critical mass of interview partners from each of the different sectors to participate in our case study. We have reasonable grounds to believe that a majority of organizations declined their participation because they do not want to disclose security-related information.

Our empirical findings (Weishäupl et al. 2017) show that (1) organizations' investments in information security are largely driven by industry-related and macro-environmental factors, such as requirements of partner firms and legal regulations. For example, the IT Security Act (ITSG) is a German federal act requires firms that operate critical infrastructures (e.g., energy, water, health or telecommunication) to improve the protection of their network against hacker attacks. (2) Firms' implementation of security processes is mainly triggered by external pressure rather than by internal incentives and attempts to reduce business risk. The reason for this is that the impact of the security processes on the business processes is judged to be negative despite its effect of increased security and expected decrease of breaches as they slow down the business processes. (3) Processes, models and methods for information security decision-making are applied rarely. Instead, decisions are made by the CISO in collaboration with the information security department (if it exists) and the CIO depending on the CISO's hierarchical position within the organization. Different opinions and preferences are discussed without using formal multi-stakeholder decision models or decision support systems. (4) Both the implementation of evaluation processes and the application of metrics hard exist. This lack of evaluation processes is rooted in the complexity and time expenditure of evaluating information security investment decisions. However, firms are urged to evaluate when external pressure exists (audits), business processes do not run smoothly, or the IT budget is reallocated. (5) Learning activities mainly occur at an ad-hoc basis and are not based on specific learning strategies. However, it was stated that for some security resources, especially for workshops in employee training, organized learning takes place because firms consider the fluctuation of employees and the fact that employees quickly forget lessons learned in past workshops.

Addressing the reviewers' comments in the intermediate evaluation report, we accounted for the multi-stakeholder aspect when they are involved in information security investment decision making.. We asked our interview partners in the case study whether there are multiple stakeholders involved in the decision-making process and if so, how conflicts due to different risk preferences are managed and solved. According to the interviewees, there are at least two stakeholders who are involved in the investment decision process. Usually, the CIO and the CISO are involved and they often express different opinions: While the CIO is mostly interested in maintaining the availability and service level of IT, the CISO's focus is on confidentiality and integrity. The CIO is made responsible if systems break down and are not available because of new security countermeasures which failed or did not work properly and were initiated by the CISO. This leads to a tradeoff discussion between CIO and CISO. It was stated that, in practice, this problem is not solved with models and data but informally discussed.

We also addressed the reviewers' suggestions to consider decision and voting processes. According to the interviews, standardized decision processes as they are known in academic literature are not implemented by organizations. In some cases, attack simulations are carried out and presented to the CFO to justify a necessary investment decision. This is kept simple: Firms use a two-dimensional matrix, either with costs-effort or cost-benefit analysis. Occasionally, a strategy pyramid is of help. Moreover, we found that voting processes are rarely conducted. Instead, different opinions among the stakeholders are discussed until a solution is found.

3.10.3 Contribution to FORSEC Research Alliance

Research project 10 is a transfer project, which is linked to the two clusters PreSTA and CLOUD. In the following, we describe the contribution of our project to each of these clusters.

This project contributed to PreSTA in terms of decision support models and techniques for preparatory defense. We developed decision models which aim at optimizing the preparedness of IT systems and support decision makers in securing their IT systems focusing on the economic perspective, i.e., considering economic constraints, such as the firm's IT security budget. Our decision models give guidance on how much to invest in which security measurement (e.g., firewall, intrusion detection, anti-virus software or security awareness training for employees) in order to attain a specified level of security to prevent security incidents or to be optimally prepared in case of incidents.

In cooperation with TP7, we proposed a decision support model to assist decision makers in organizations in the challenging task whether to invest in an Identity and Access Management System (IAMS) and if so, in which kind, e.g., implementing role-based access control or access control lists (Weishäupl, Kunz, et al. 2015). We quantified the benefits of security and cost savings of the implementation of various types of IAMS and took the firm's risk preferences into account. This model can be used by organizations to make optimal decisions regarding both security and costs.

Our project contributed to cluster CLOUD by taking into account economic aspects in the recovery, auditing and forensics phase. The main contribution was the provision of measurements of the financial loss of downtime of virtual machines and the quantification of financial aspects of non-compliance with QoS agreements. We provided (1) a quantitative analysis of reputational damages and financial losses as a consequence of downtimes and (2) the quantification of financial damages of non-compliance with quality of service agreements.

In various joint publications with the cluster and with TP4, we derived the following conclusions (Fischer et al. 2015; Mandarawi and Weishäupl 2017; Mandarawi et al. 2015; Rakotondravony, Taubmann, et al. 2017): Although cloud computing offers the possibility of cost savings for cloud users through optimized and efficient computing, the benefits are accompanied by numerous security treats and risks that can lead to financial harm for the cloud provider. In the case of an attack not only the costs for working hours for analyzing, repairing and disinfecting the systems and losses in productivity, revenue and reputation need to be considered. The cloud provider also needs to pay penalties to the cloud user when the quality of service specified in the service level agreements is not given.

3.10.4 Beyond FORSEC

As our integrated view on information security investments covers multi-disciplinary effects, including economic, technologic and regulatory issues, our results provide a comprehensive theoretical foundation for future research in the field of management of IT security investments. In particular, we plan to use the developed theory, models and metrics in cooperation with firms for further empirical research with a focus on embedding them into real-world processes, implementing decision support systems and acquiring security-related data to support decision making and learning. We intend to further collaborate with the organizations that already participated in our multiple case-study.

From a theoretical point of view, the learning behavior of strategic attackers has been modelled in the literature with game theoretical approaches. We intend to draw on these concepts to determine an optimal information security investment strategy by considering attackers' adaptive behavior. This is of particular importance as attackers learn from their past errors and find new ways to exploit vulnerabilities so that, in turn, firms need to adapt their own behavior and countermeasures. Considering games with incomplete information, we take into account that both the firm and the attacker face informational uncertainty.

3.10.5 Publications

Rauchecker et al. (2014): A Decision Support System for IT Security Incident Management	
Abstract	The problem of processing IT security incidents is a key task in the field of security service management. This paper addresses the problem of effectively assigning and scheduling security incidents to the members of the IT staff. To solve this problem, we propose an innovative approach to assign staff members to security incidents by applying mathematical programming to the field of IT security management. We formulate an optimization model and propose efficient solution methods. The numerical simulations show that our approach improves current best practice behaviour significantly.
Citation	Rauchecker, G., Yasasin, E., and Schryen, G. 2014. “A Decision Support System for IT Security Incident Management,” in *Proceedings of the 11th International Conference on Trust, Privacy and Security in Digital Business (TRUSTBUS 2014)*, pp. 36–47.
URL	https://link.springer.com/chapter/10.1007/978-3-319-09770-1_4

Weishäupl, Kunz, et al. (2015): Towards an Economic Approach to Identity and Access Management Systems Using Decision Theory	
Abstract	Nowadays, providing employees with failure-free access to various systems, applications and services is a crucial factor for organizations' success as disturbances potentially inhibit smooth workflows and thereby harm productivity. However, it is a challenging task to assign access rights to employees' accounts within a satisfying time frame. In addition, the management of multiple accounts and identities can be very onerous and time consuming for the responsible administrator and therefore expensive for the organization. In order to meet these challenges, firms decide to invest in introducing an Identity and Access Management System (IAMS) that supports the organization by using policies to assign permissions to accounts, groups, and roles. In practice, since various versions of IAMSs exist, it is a challenging task to decide upon introduction of an IAMS. The following study proposes a first attempt of a decision support model for practitioners which considers four alternatives: Introduction of an IAMS with Role-based Access Control (RBAC) or without and no introduction of IAMS again with or without RBAC. To underpin the practical applicability of the proposed model, we parametrize and operationalize it based on a real world use case using input from an expert interview.
Citation	Weishäupl, E., Kunz, M., Yasasin, E., Wagner, G., Prester, J., Schryen, G., and Pernul, G. 2015. “Towards an Economic Approach to Identity and Access Management Systems Using Decision Theory,” in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015),* Vienna, Austria.

Weishäupl, Yasasin, et al. (2015b): IT Security Investments through the Lens of the Resource-based View: A New Theoretical Model and Literature Review	
Abstract	IT security has become a major issue for organizations as they need to protect their assets, including IT resources, intellectual property and business processes, against security attacks. Disruptions of IT-based business activities can easily lead to economic damage, such as loss of productivity, revenue and reputation. Organizations need to decide (1) which assets need which level of protection, (2) which technical, managerial and organizational security countermeasures lead to this protection and (3) how much should be spent on which countermeasure in the presence of budget constraints. Answering these questions requires both making IT security investment decisions and evaluating the effectiveness and efficiency of these decisions. The literature has contributed to this field adopting approaches from micro-economics, finance and management, among others. However, the literature is rather fragmented and lacks a shared theoretical basis. As a consequence, it remains partly open what we can learn from past research and how we can direct and stimulate still missing research activities. In order to address these deficiencies, we draw on the resource-based view (RBV) and provide a theoretical model for IT security investments. We use this RBV model to review the IT security investment literature and to identify research gaps.
Citation	Weishäupl, E., Yasasin, E., and Schryen, G. 2015b. “IT Security Investments through the Lens of the Resource-based View: A new theoretical Model and Literature Review,” in *Proceedings of the 23rd European Conference on Information Systems (ECIS 2015),* Münster, Germany.
URL	http://aisel.aisnet.org/ecis2015_cr/198/

Weishäupl, Yasasin, et al. (2015a): A Multi-Theoretical Literature Review on Information Security Investments using the Resource-Based View and the Organizational Learning Theory	
Abstract	The protection of information technology (IT) has become and is predicted to remain a key economic challenge for organizations. While research on IT security investment is fast growing, it lacks a theoretical basis for structuring research, explaining economic-technological phenomena and guide future research. We address this shortcoming by suggesting a new theoretical model emerging from a multi-theoretical perspective adopting the Resource-Based View and the Organizational Learning Theory. The joint application of these theories allows to conceptualize in one theoretical model the organizational learning effects that occur when the protection of organizational resources through IT security countermeasures develops over time. We use this model of IT security investments to synthesize findings of a large body of literature and to derive research gaps. We also discuss managerial implications of (closing) these gaps by providing practical examples.
Citation	Weishäupl, E., Yasasin, E., and Schryen, G. 2015a. “A Multi-Theoretical Literature Review on Information Security Investments using the Resource-Based View and the Organizational Learning Theory,” in *Proceedings of the 36th International Conference of Information Systems (ICIS 2015)*, Fort Worth, TX.
URL	http://aisel.aisnet.org/icis2015/proceedings/SecurityIS/16/

Weishäupl (2017): Towards a Multi-objective Optimization-model to Support Information Security Investment Decision-making	
Abstract	The protection of assets, including IT resources, intellectual property and business processes, against security attacks has become a challenging task for organizations. From an economic perspective, firms need to minimize the probability of a successful security incident or attack while staying within the boundaries of their information security budget in order to optimize their investment strategy. In this paper, an optimization model to support information security investment decision-making in organizations is proposed considering the two conflicting objectives (simultaneously minimizing the costs of countermeasures while maximizing the security level). Decision models that support the firms' decisions considering the trade-obetween the security level and the investment allocation are beneficial for organizations to facilitate and justify security investment choices.
Citation	Weishäupl, E. 2017. "Towards a Multi-objective Optimization-model to Support Information Security Investment Decision-making," in *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems (SHCIS 2017)*, Neuchâtel, Switzerland.

Yasasin et al. (2014): A Fuzzy Security Investment Decision Support Model for Highly Distributed Systems	
Abstract	The economic aspect of information security is a comparatively new discipline so that there is hardly any extensive research work. This applies in particular to measures in highly distributed systems which have been neglected in previous research. The present paper focuses on the security investments in such systems. We augment an existing research about a fuzzy decision support model by defining appropriate operators in order to applicate this work in practice. The proposed model includes uncertainty with respect to the impact of investments on the achieved security levels of components of the distributed system. We further develop a heuristic to solve the problem and test the heuristic experimentally. The paper concludes with a discussion and gives an outlook to future work in the context of security investments.
Citation	Yasasin, E., Rauchecker, G., Prester, J., and Schryen, G. 2014. “A Fuzzy Security Investment Decision Support Model for Highly Distributed Systems,” in *Proceedings of the 1st Workshop on Security in Highly Connected IT Systems (SHCIS 2014),* Munich, Germany.
URL	http://ieeexplore.ieee.org/document/6974864/

Yasasin and Schryen (2015): Requirements for IT security Metrics - An Argumentation Theory Based Approach	
Abstract	The demand for measuring IT security performance is driven by regulatory, financial, and organizational factors. While several best practice metrics have been suggested, we observe a lack of consistent requirements against which IT security metrics can be evaluated. We address this research gap by adopting a methodological approach that is based on argumentation theory and an accompanying literature review. As a result, we derive five key requirements: IT security metrics should be (a) bounded, (b) metrically scaled, (c) reliable, valid and objective, (d) context-specific and (e) computed automatically. We illustrate and discuss the context-specific instantiation of requirements by using the practically used "vulnerability scanning coverage" and "mean-time-to-incident discovery" metrics as examples. Finally we summarize further implications of each requirement.
Citation	Yasasin, E., and Schryen, G. 2015. “Requirements for IT security Metrics - An Argumentation Theory Based Approach,” in *Proceedings of the 23rd European Conference on Information Systems (ECIS 2015),* Münster, Germany.
URL	http://aisel.aisnet.org/ecis2015_cr/208/

3.11 TP11 – Security and Data Protection in Smart Grid

3.11.1 Project Overview

The energy system is in a rapid change and becomes more and more connected. New IoT technology reaches the households and the customer's Home Area Network (HAN) communicates over the grid structure with different entrepreneurs. This new critical infrastructure is called Smart Grid and offers a bunch of new attack vectors.

The overall focus of TP11 is the security and data protection in Smart Grid. The specific goals of the project was in a first step to analyse the kind and amount of information flow from sensor data to third parties. To protect the prosumers from malware and attackers, a privacy-enhanced decentralized intrusion and fraud detection concept was developed. A neuronal network therefore analyses the overall consumption of the household, to recognise known appliances and detect irregularities. All consumption data is processed inside the Home Area Network. This means that no privacy-critical data has to be transferred.

3.11.2 Results Achieved

The first task of work package (WP1), was to conduct a literature and requirement study on the state of the art. The results were published in (Richthammer 2014; Scheibel 2014). We have discovered that the topic Smart Grid is well analyzed, but publications with the combination Smart Grid and privacy are very rare. To measure the impact of single publications

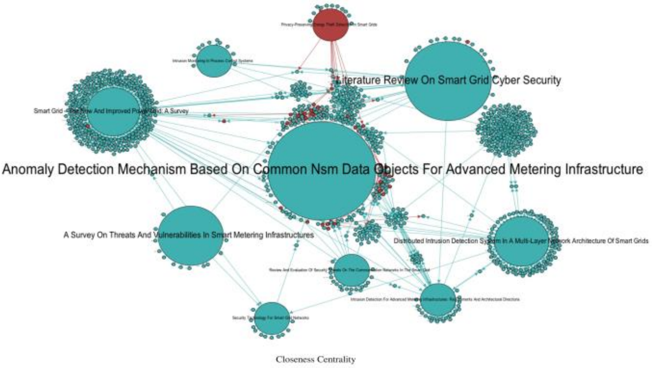
Privacy-Preserving Energy Theft Detection In Smart Grids
Intrusion Monitoring In Process Control Systems
Literature Review On Smart Grid Cyber Security
Smart Grid – The New And Improved Power Grid: A Survey
Anomaly Detection Mechanism Based On Common Nsm Data Objects For Advanced Metering Infrastructure
A Survey On Threats And Vulnerabilities In Smart Metering Infrastructures
Distributed Intrusion Detection System In A Multi-Layer Network Architecture Of Smart Grids
Review And Evaluation Of Security Threats On The Communication Networks In The Smart Grid
Intrusion Detection For Advanced Metering Infrastructures: Requirements And Architectural Directions
Security Technology For Smart Grid Networks
Closeness Centrality

Figure 6: Closeness Centrality graph

we generated a network with directed graphs that shows the interconnections between different papers and their references. Figure 6 depicts the results of the analysis. We used the "Closeness Centrality" metric, which measures how far a node is away from others. A high value means that the publication and their references use lot of third party references. This suggests that these are survey papers which is confirmed by a closer look.

The node size visualises the closeness centrality value. The red coloured nodes represent publications that are related to privacy. As one can see, there is just one node of relevance, which is related to the privacy topic. This affirm our assumption that privacy is not yet sufficiently investigated. A closer look to the publications validate the assumption. The most considered publication does not address privacy at all or if so in only a small amount.

WP2 addresses the information flow problem of Smart Grid and was published in (Richthammer 2014; Scheibel 2014). Abstract representations of information flows and measures have been addressed in paper (Richthammer and Reif 2015). More concrete information flow analyses were considered with respect to sensor data in (Scheibel 2014). The analysis is based on the BSI Protection Profile, which has a strong influence on privacy legislations in the European Union. The evaluators also suggested to consider privacy and data protection, which was also analyzed as part of (Richthammer 2014; Scheibel 2014). Depending on the tariff, fine-grained energy consumption data is transferred to the Energy Service Provider and the Grid Operator. The default capture interval is set to every 15 minutes but can be downsized up to every minute. With such a fine-grained consumption trace the daily routine and behaviour prediction of house inhabitants is possible. For example from the information when people get up, up to the television consumption behaviour which channel is watched.

Paper (Richthammer 2014) contributes to WP2 and WP3. Paper (Richthammer 2014) discusses the privacy issue for pseudonymised customers in the Smart Grid infrastructure, while paper (Pham and Kesdogan 2015) addresses anonymity problems that arise from information disclosure. Papers (Pham and Kesdogan 2015; Richthammer 2014) offer a technical perspective on a potential attack on a Smart Grid Infrastructure. Based on this knowledge an approach for a secure and privacy focused device management for customers was developed and published in (Richthammer and Reif 2015). These findings will be used in WP3 and WP4. The idea behind this concept is a privacy-friendly Intrusion and Fraud Detection System. The developed IDS architecture is depicted in Figure 7 and structured in three parts. The Information Collection Module gathers the energy consumption from the Smart Meter. In addition, metadata such as meteorological information like the temperature, humidity, hydrostatic pressure and lux are collected. The Information Collecting Policy defines, how this data is collected. The collected information results in Consumption Events and is processed with the help of the NILMTK Framework, created by Nipun Batra, Jack Kelly and Oliver Parson in 2014.

The Detection Module is the core component of the Intrusion and Fraud Detection System. It processes the Consumption Events and tries to detect anomalies with a neuronal network algorithm. Based on the Detection Policy and the System Information, the IDS Sensor decides if the actual consumption is irregular and tries to detect malicious or failed appliances. The gathered Consumption Events are analyzed using a Recurrent Neural Network (RRN). In our work we used a Long short-term memory (LSTM) algorithm implemented in python.

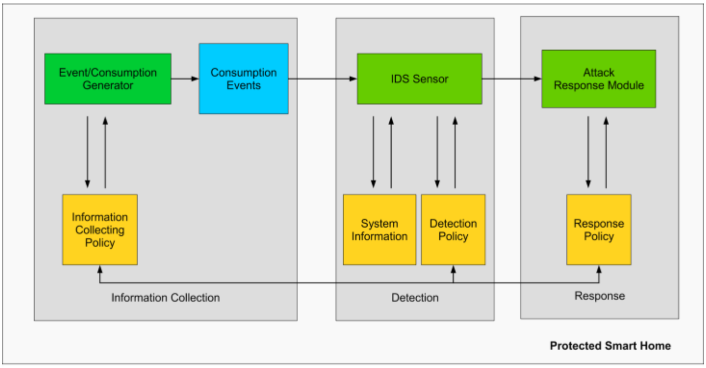

Figure 7: Intrusion Detection System structure

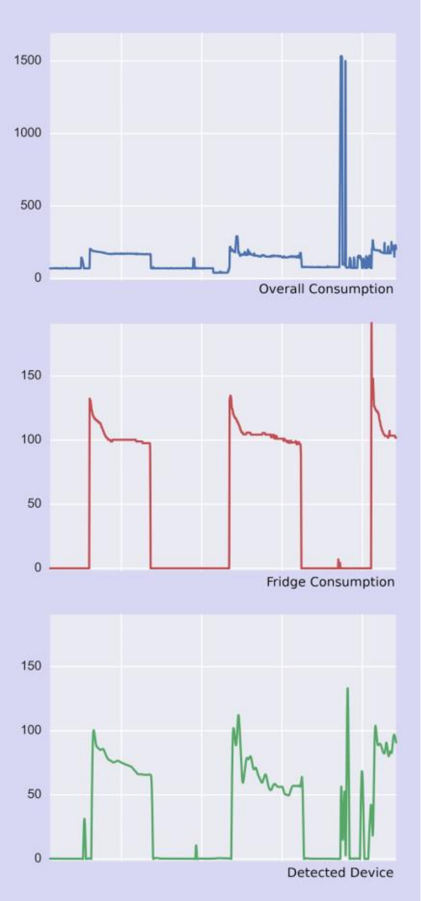

Figure 8: Evaluation of the algorithm

A good implementation of the RRN algorithms is realised by the Keras framework, which is an actively maintained open-source project. We used the Keras framework as a core basis for our research and optimised the LSTM training and detection parameters. During the training phase our LSTM algorithm uses a time frame of 512 datapoints of the overall consumption as input neurons. The last datapoint of the 512 time frame from every appliance was used as output neuron. For every training iteration the time frame was shifted by one unit. Every device was trained separately. Hence it is possible to detect malicious infections individually and more precisely for every appliance within a household. A System Information module provides the IDS Sensor with additional information and data. Meteorological information for example can support the decision process, if an irregularity happens due to an attack or environmental influences. For example high consumption based on an air condition system can be caused by an attack or an unusual hot day.

The Response Module defines, how a detected irregularity is processed. A Response Policy provides this information and can be updated and extended by the prosumer. For example the user is able to configure how to be be informed via e-mail, sms or push information in combination with a smartphone application. Consumption information can be transferred to the application. How this data can be handled and processed securely is considered by TP1 and TP2.

WP5 addresses the evaluation of the developed concept. Figure 8 shows one result of the evaluation of our algorithm. The blue line is the overall consumption of an household. The red line stand for the real consumption of an appliance inside the household. In our example we evaluated the algorithm for a refrigerator. In the overall consumption the energy consuming fridge is visually recognizable. Every consumption peak shown in the red line leads to an increased overall consumption. The last consumption picture with the green line shows the predict of our detection algorithm. The LTSM based software sensor goes over the overall consumption and tries to predict the activation time of the refrigerator. As can be seen, the predicted green line overlaps nearly exactly the real red consumption of the fridge. With a normalised trace it is now possible to detect irregularities of an appliance. As well in the consumption intensity as also in the consumption intervall. The final results are planned to be published at a conference for neuronal networks with a focus on security. A decision for the conference has not been made until the deadline of the final report.

3.11.3 Contribution to FORSEC Research Alliance

Our main contribution to PreSTA was the developed concept from WP3 and WP4. This concept was the basis for the research of other TPs on usability and security awareness, economic aspects, reputation systems and identity management systems. We were involved in many discussions at workshops and provided input for other TPs.

We had cooperations with the the colleagues from TP8. The concepts of reputation systems, the main focus of TP8, was used for some parts of the visual analyses of events supported by graph generation. The group from TP8 examines purchase events and tries to detect fraud and selfish behaviour of market place seller. These are also time-dependent and connected events, similar to energy consumption events. Here we had many fruitful knowledge exchanges with the working group of TP8. Especially regarding ways to visualize big data events.

TP10 considered the economical aspects within IT-security concepts. The economic aspects of Smart Grids in combination with privacy aspects was analyzed. To detect security-relevant incidents and malicious devices. A feasibel algorithm is necessary. But also the economic aspect of a developed concept has to be considered. The best solution will not be used if it is too expensive to implement. Fruitful knowledge transfer took place during the project period and ended up in an analysis of economic investments in the field of IT-Security. Details will be discussed in the chapter of TP10.

In WP4 we implemented our concept in a working algorithm. For WP5 the evaluation and usability analysis of a developed solution was planned. Our workshops with our colleagues from TP6 lead to new conclusions about the awareness of users in relation to security solutions. The evaluations from TP6 have shown, that usability of IT-security products is an important factor for users. Until the deadline of the final report a usability friendly Graphical User Interfaces was not finished but will be until the project ending. Therefore, a final knowledge exchange will take place with TP6.

3.11.4 Beyond FORSEC

Before we began developing our concept of a decentralised Intrusion Detection System, we started with a literature review and state of the art analysis. As it turned out, the idea of appliances detection based on the overall energy consumption is well researched. The first publication in this field was from Hart 1992. Since then there have been a lot of publications in this field. Many algorithms and new concepts have been developed and improved. What most of all implementations have in common, the detection rate under laboratory conditions is throughout good, but under real conditions the detection rate breaks down.

In the past, there were some conceptual ideas to combinate NILM with Neuronal Network technologies. But due to the fact that only CPU computation was possible, the training phase took very long. Nowadays we have much more powerful computers, especially in combination with GPU computing. This was the reason why we picked up the idea and developed an concept, based on the state of the art.

With our project we have demonstrated that Neuronal Network learning with GPU power was much more efficient, thus opening the door for research in the area of NILM appliance detection.

3.11.5 Publications

Richthammer and Reif (2015): Intrusion Detection in the Smart Grid based on an Analogue Technique	
Abstract	In Smart Grid a customer’s privacy is threatened by the fact that an attacker could deduce personal habits from the detailed consumption data. We analysed the publications in this field of research and found out that privacy does not seem to be the main focus. To verify this guess, we analysed it with the technique of directed graphs. This indicates that privacy isn’t yet sufficiently investigated in the Smart Grid context. Hence we suggest a decentralised IDS based on NILM technology to protect customer’s privacy. Thereby we would like to initiate a discussion about this idea.
Citation	Richthammer, H., and Reif, S. 2015. “Intrusion Detection in the Smart Grid based on an Analogue Technique,” in International Workshop on Open Problems in Network Security (iNetSec 2015).
URL	https://link.springer.com/chapter/10.1007/978-3-319-39028-4_5

Richthammer and Kesdogan (2015): Secure and Privacy Focused Customer Device Management in a Smart Household Environment	
Abstract	In our modern world the energy management has to change, because the traditional roles between the energy producers and the energy consumers are already in a transformation process. One solution could be the Smart Grid. The Energy Service Provider and the consumer are connected over a bidirectional network. Therefor consumption (and other digital) information can be exchanged easily between the participants. To provide such a solution a `smart' connected household with an implemented Smart Meter is necessary. The Smart Meter is the connection between the household and the outer world. In the future we will get a lot of `smart' devices which can interact with the Smart Meter. But there will also be old devices with no `smart' technology. In our approach we will show how the customer is able to manage all, also the `not smart' devices and protection his privacy at the same time in an adequate way.
Citation	Richthammer, H., and Kesdogan, D. 2015. "Secure and Privacy Focused Customer Device Management in a Smart Household Environment," in *Proceedings of the 30th IFIP International Information Security and Privacy Conference (IFIP SEC 2015).*

Richthammer (2014): The Privacy Issues for Pseudonymised Customers in the Smart Grid	
Abstract	This is a short overview of the privacy issues for customers in the Smart Grid infrastructure. Smart meter and other devices within the Smart Grid produce a lot of privacy sensitive data. With this find grained data it is possible to make predictions about the daily routine of a household or create a movement profile of a vehicle. There are techniques to protect the privacy of a person in network-alike structures, e.g. by creating pseudonyms or anonymizing the flow of data. But this protection does not always work in an adequate way, for example if there are quasi identifier, significant or linked pattern, it is possible to depseudonymize and de-anonymize a customer. So it is possible to analyse the daily routine of a person as well as creating a movement profile of his electrical vehicle or getting some informations about his preferences or issues.
Citation	Richthammer, H. 2014. “The Privacy Issues for Pseudonymised Customers in the Smart Grid,” in *Proceedings of the International Symposium on Engineering Secure Software and Systems (ESSoS 2014)*.

Pham and Kesdogan (2015): Towards Relations between the Hitting-Set Attack and the Statistical Disclosure Attack	
Abstract	The Minimal-Hitting-Set attack (HS-attack) is a well-known, provably optimal exact attack against the anonymity provided by Chaumian Mixes (Threshold-Mixes). This attack allows an attacker to identify the fixed set of communication partners of a given user by observing all messages sent and received by a Chaum Mix. In contrast to this, the Statistical Disclosure attack (SDA) provides a guess of that user's contacts, based on statistical analyses of the observed message exchanges. We contribute the first closed formula that shows the influence of traffic distributions on the least number of observations of the Mix to complete the HS-attack. This measures when the Mix fails to hide a user's partners, such that the user cannot plausibly deny the identified contacts. It reveals that the HS-attack requires asymptotically less observations to identify a user's partners than the SDA, which guesses them with a given bias. This number of observations is O($\frac{1}{p}$) for the HS-attack and O($\frac{1}{p^2}$) for the SDA, where p the probability that the attacked user contacts his least frequent partner.
Citation	Pham, D. V., and Kesdogan, D. 2015. "Towards Relations between the Hitting-Set Attack and the Statistical Disclosure Attack," in *Proceedings of the 30th IFIP International Information Security and Privacy Conference (IFIP SEC 2015)*, Hamburg, Germany.
URL	https://link.springer.com/chapter/10.1007/978-3-319-18467-8_3

Holler et al. (2015): Practical Attacks on pxe based Boot Systems	
Abstract	Sustainable energies are an important resource for the twenty-first century and the generation of renewable energy raised during the last years. Decentralised energy production rises with households being able to feed wind and solar power into the grid. To handle this bidirectional flow of energy and hold the whole system in balance, an innovative management is necessary. This was the birth of the smart grid idea. The organisation and handling of the grid and all stakeholders is managed in a Supervisory Control And Data Acquisition (SCADA) system centre. To administrate all computer devices within a SCADA centre, one solution might be thin clients with Preboot Execution Environment (PXE) boot up because network boot systems are an elementary part of a modern system management solution. An important fact of such an approach is the security aspect. However, the PXE standard comes without any security functionality. In our research, we analyse possible attack vectors and create a testbed to simulate different attacks on such an infrastructure successfully. The attack vectors are based on the protocols Dynamic Host Configuration Protocol (DHCP) and Trivial File Transfer Protocol (TFTP).
Citation	Holler, P., Roth, C., and Richthammer, H. 2015. “Practical Attacks on PXE based Boot Systems,” in *Proceedings of the 1st International Conference on Cyber Security for Sustainable Society (CSSS 2015)*, Coventry, UK.

4. Publications

Bauer, R., Staudemeyer, C., Pöhls, H. C., and Fragkiadakis, A. 2016. "ECDSA on Things: IoT Integrity Protection in Practise," in *Proceedings of the International Conference on Information and Communication Systems (ICICS 2016)*, Irbid, Jordan.

Benenson, Z., Freiling, F., Glass, B., Hänsch, N., Müller, T., Protsenko, M., and Zhuang, Y. 2017. "Software Obfuscation Causes Experienced Users to Reverse Engineer Like Beginners," *(in preparation)*.

Benenson, Z., Gassmann, F., and Landwirth, R. 2017. "Unpacking Spear Phishing Susceptibility," in *Proceedings of the 21st International Conference on Financial Cryptography and Data Security*, Malta.

Benenson, Z., Girard, A., Hintz, N., and Luder, A. 2014. "Susceptibility to URL-based Internet Attacks: Facebook vs. email," in *Proceedings of the 6th IEEE International Workshop on SEcurity and SOCial Networking (SESOC 2014)*, Budapest, Hungary.

Benenson, Z., Girard, A., and Krontiris, I. 2015. "User Acceptance Factors for Anonymous Credentials: An Empirical Investigation," in *Proceedings of the 14th Workshop on the Economics of Information Security (WEIS 2015)*, Delft, Netherlands.

Benenson, Z., Girard, A., Krontiris, I., Liagkou, V., Rannenberg, V., and Stamatiou, Y. 2014. "User acceptance of Privacy-ABCs: An Exploratory Study," in *Proceedings of the 2nd International Conference on Human Aspects of Information Security, Privacy, and Trust*, Crete, Greece.

Benenson, Z., Lenzini, G., Oliveira, D., Parkin, S., and Uebelacker, S. 2015. "Maybe Poor Johnny Really Cannot Encrypt - The Case for a Complexity Theory for Usable Security," in *New Security Paradigms Workshop*, Twente, Netherlands.

Böhm, A. 2017. "Live Migration of two Linked Machines in Xen (in progress),", University of Passau.

Braun, B., Gries, C., Petschkuhn, B., and Posegga, J. 2014. "Ghostrail: Ad Hoc Control-Flow Integrity for Web Applications," in *Proceedings of the 29th IFIP International Information Security and Privacy Conference (IFIP SEC 2014)*.

Braun, B., Köstler, J., Posegga, J., and Johns, M. 2014. "A Trusted UI for the Mobile Web," in *Proceedings of the 29th IFIP International Information Security and Privacy Conference (IFIP SEC 2014)*.

Braun, B., Pauli, K., Posegga, J., and Johns, M. 2015. "LogSec: Adaptive Protection for the Wild Wild Web," in *2015 ACM Symposium on Applied Computing (SAC 2015)*.

Busch, M., Protsenko, M., and Müller, T. 2015. "Automated Malware Analysis for Android: A Comparative Evaluation," in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015)*, Vienna, Austria.

Cetto, A., Netter, M., Pernul, G., Richthammer, C., Riesner, M., Roth, C., and Sänger, J. 2014. "Friend Inspector: A Serious Game to Enhance Privacy Awareness in Social Networks," in *Proceedings of the 2nd International Workshop on Intelligent Games for Empowerment and Inclusion (IDGEI 2014)*, Haifa, Israel.

Dawaras, S. 2015. "Theoretische und Praktische Untersuchung von Trusted-Computing-Techniken für eine IaaS-Cloud,", University of Passau.

Dresel, L., Protsenko, M., and Müller, T. 2016. "ARTIST: The Android Runtime Instrumentation Toolkit," in *Proceedings of the 11th International Conference on Availability, Reliability and Security (ARES 2016)*, Salzburg, Austria.

Fischer, A., and de Meer, H. 2016. "Generating Virtual Network Embedding Problems with Guaranteed Solutions," *IEEE Transactions on Network and Service Management* (13:3), pp. 504–517.

Fischer, A., Kittel, T., Kolosnjaji, B., Lengyel, T. K., Mandarawi, W., Reiser, H. P., Taubmann, B., Weishäupl, E., de Meer, H., Mu, T., and Protsenko, M. 2015. "CloudIDEA: A Malware Defense Architecture for Cloud Data Centers," in *Proceedings of the 5th International Symposium on Cloud Computing, Trusted Computing and Secure Virtual Infrastructures - Cloud and Trusted Computing (C&TC)*, Rhodes, Greece.

Fischer, A., Kuehn, R., Mandarawi, W., and de Meer, H. 2016. "Modeling Security Requirements for VNE algorithms," in *Proceedings of the 10th EAI International Conference on Performance Evaluation Methodologies and Tools*.

Frädrich, C., Pöhls, H. C., Popp, W., Rakotondravony, N., and Samelin, K. 2016. "Integrity and Authenticity Protection with Selective Disclosure Control in the Cloud and IoT," in *Proceedings of the International Conference on Information and Communication Systems (ICICS 2016),* Irbid, Jordan.

Freiling, F., Glanzmann, T., and Reiser, H. P. 2017. "Characterizing loss of digital evidence due to abstraction layers," *Digital Investigation* (20), pp. 107–115.

Freiling, F., Protsenko, M., and Zhuang, Y. 2014. "An Empirical Evaluation of Software Obfuscation Techniques applied to Android APKs," in *Proceedings of the International Workshop on Data Protection in Mobile and Pervasive Computing (DAPRO 2014)*, Beijing, China.

Fuchs, L., Kunz, M., and Pernul, G. 2014. "Role Model Optimization for Secure Role-based Identity Management," in *Proceedings of the 22nd European Conference on Information Systems (ECIS 2014)*, Tel Aviv, Israel.

Guerra, M. C. 2017. "Dynamic Security Analysis using Virtual Machine Introspection (in progress),", Instituto Superior Técnico, Lisboa.

Hahn, S., Protsenko, M., and Müller, T. 2016. "Comparative Evaluation of Machine Learning-based Malware Detection on Android," in *Sicherheit 2016*.

Hänsch, N., and Benenson, Z. 2014. "Specifying IT Security Awareness," in *Proceedings of the 1st Workshop on Security in Highly Connected IT systems (SHCIS 2014),* Munich, Germany.

Hassan, S., Sänger, J., and Pernul, G. 2014. "SoDA: Dynamic Visual Analytics of Big Social Data," in *Proceedings of the 1st International Conference on Big Data and Smart Computing (BigComp 2014)*, Bangkok, Thailand.

Haupert, V., and Müller, T. 2016. "Auf dem Weg verTAN: Über die Sicherheit App-basierter TAN-Verfahren," in *Sicherheit 2016*.

Hintz, N., Engelberth, M., Benenson, Z., Freiling, F., Hintz, N., Engelberth, M., Benenson, Z., and Freiling, F. 2014. "Phishing still works: Erfahrungen und Lehren aus der Durchführung von Phishing-Experimenten," in *GI-Sicherheit 2014*.

Holler, P., Roth, C., and Richthammer, H. 2015. “Practical Attacks on PXE based Boot Systems,” in *Proceedings of the 1st International Conference on Cyber Security for Sustainable Society (CSSS 2015),* Coventry, UK.

Huber, M., Taubmann, B., Wessel, S., Reiser, H. P., and Sigl, G. 2016. “A Flexible Framework for Mobile Device Forensics Based on Cold Boot Attacks,” *EURASIP Journal on Information Security*.

Hummer, M., Kunz, M., Fuchs, L., and Pernul, G. 2015. “Advanced Identity and Access Policy Management using Contextual Data,” in *Proceedings of the 10th International Conference on Availability, Reliability and Security (ARES 2015)*, Toulouse, France.

Hummer, M., Kunz, M., Fuchs, L., and Pernul, G. 2016. “Adaptive Identity and Access Management - Contextual Data based Policies,” *EURASIP Journal on Information Security* (19).

Huppert, P. 2015. “Virtual Machine Introspection During Live Migration,”, University of Passau.

Kolosnjaji, B., and Eckert, C. 2015a. “Neural Network-Based User-Independent Physical Activity Recognition for Mobile Devices,” in *Proceedings of the 16th Conference on Intelligent Data Engineering and Automated Learning (IDEAL 2015)*, Wroclaw, Poland: Springer, pp. 378–386.

Kolosnjaji, B., and Eckert, C. 2015b. “Leveraging Deep Learning for Malware Detection and Classification,” in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015),* Vienna, Austria.

Kolosnjaji, B., Eraisha, G., Webster, G., Zarras, A., and Eckert, C. 2017. “Empowering Convolutional Networks for Malware Classification and Analysis,” in *Proceedings of the 30th International Joint Conference on Neural Networks (IJCNN 2017)*, Anchorage, AK.

Kolosnjaji, B., Zarras, A., Lengyel, T., Webster, G., and Eckert, C. 2016. “Adaptive Semantics-Aware Malware Classification,” in *Proceedings of the 13th Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA 2016),* San SebastiÃ¡n, Spain.

Krontiris, I., Benenson, Z., Gerard, A., Sabouri, A., Rannenberg, K., and Schoo, P. 2015. “Privacy-ABCs as a Case for Studying the

Adoption of PETs by Users and Service Providers," in *Annual Privacy Forum*.

Kunz, M., Fuchs, L., Hummer, M., and Pernul, G. 2015. "Introducing Dynamic Identity and Access Management in Organizations," in *Proceedings of the 11th International Conference on Information Systems Security (ICISS 2015),* Kolkata, India.

Kunz, M., Fuchs, L., Netter, M., and Pernul, G. 2015a. "Analyzing Quality Criteria in Role-based Identity and Access Management," in *Proceedings of the 1st International Conference on Information Systems Security and Privacy (ICISSP 2015),* Angers, France.

Kunz, M., Fuchs, L., Netter, M., and Pernul, G. 2015b. "How to Discover High-quality Roles? A Survey and Dependency Analysis of Quality Criteria in Role Mining," *Communications in Computer and Information Science* (596).

Kunz, M., Hummer, M., Fuchs, L., Netter, M., and Pernul, G. 2014. "Analyzing Recent Trends in Enterprise Identity Management," in *Proceedings of the 1st Workshop on Security in Highly Connected IT Systems (SHCIS 2014),* Munich, Germany.

Lengyel, T., Kittel, T., and Eckert, C. 2015. "Virtual Machine Introspection with Xen on ARM," in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015),* Vienna, Austria.

Lengyel, T., Kittel, T., Webster, G., and Torrey, J. 2014. "Pitfalls of Virtual Machine Introspection on Modern Hardware," in *Proceedings of the 1st Workshop on Malware Memory Forensics (MMF 2014),* New Orleans, LA, December (available at https://www.sec.in.tum.de/assets/Uploads/pitfalls-virtual-machine.pdf).

Lengyel, T., Maresca, S., Payne, B. D., Webster, G. D., Vogl, S., and Kiayias, A. 2014. "Scalability, Fidelity and Stealth in the DRAKVUF Dynamic Malware Analysis System," in *Proceedings of the 30th Annual Computer Security Applications Conference (ACSAC2014),* New Orleans, LA.

Maier, D., Müller, T., and Protsenko, M. 2014. "Divide-and-Conquer: Why Android Malware cannot be stopped," in *Proceedings of the 9th International Conference on Availability, Reliability and Security (ARES 2014),* Fribourg, Switzerland.

Maier, D., Protsenko, M., and Müller, T. 2015. "A Game of Droid and Mouse: The Threat of Split-Personality Malware on Android," in *Computers & Security (COSE)*.

Mandarawi, W., Fischer, A., de Meer, H., and Weishäupl, E. 2015. "QoS-Aware Secure Live Migration of Virtual Machines," in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015),* Vienna, Austria.

Mandarawi, W., Fischer, A., Houyou, A. M., Huth, H.-P., and de Meer, H. 2016. "Constraint-Based Virtualization of Industrial Networks," in *Principles of Performance and Reliability Modeling and Evaluation*, pp. 567–586.

Mandarawi, W., and Weishäupl, E. 2017. "Performance, Privacy and Security Virtual Network SLAs and Deployment Policies in IaaS Clouds," in *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems (SHCIS 2017),* Neuchâtel, Switzerland.

Marktscheffel, T., Gottschlich, W., Popp, W., Werli, P., Fink, S. D., Bilzhause, A., and de Meer, H. 2016. "QR Code Based Mutual Authentication Protocol for Internet of Things," in *Proceedings of the 5th Workshop on IoT-SoS: Internet of Things Smart Objects and Services (WOWMOM SOS-IOT 2016).*

Naumann, J., Protsenko, M., and Müller, T. 2015. "Google Verify Apps: The Illusion of Security?," in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015),* Vienna, Austria.

Parra, J. D., Schreckling, D., and Posegga, J. 2016. "Addressing Data-Centric Security Requirements for IoT-Based Systems," in *Proceedings of the of International Workshop on Secure Internet of Things (SIOT 2016),* Oslo, Norway.

Pham, D. V., and Kesdogan, D. 2015. "Towards Relations between the Hitting-Set Attack and the Statistical Disclosure Attack," in *Proceedings of the 30th IFIP International Information Security and Privacy Conference (IFIP SEC 2015),* Hamburg, Germany.

Protsenko, M., Kreuter, S., and Müller, T. 2015. "Dynamic Self-Protection and Tamperproofing for Android Apps using Native Code," in *Proceedings of the 10th International Conference on Availability, Reliability and Security (ARES 2015),* Toulouse, France.

Protsenko, M., and Müller, T. 2013. “PANDORA Applies Non-Deterministic Obfuscation Randomly to Android,” in *Proceedings of the 8th International Conference on Malicious and Unwanted Software (Malware 2013)*, Fajardo, Puerteo Rico.

Protsenko, M., and Müller, T. 2014. “Android Malware Detection based on Software Complexity Metrics,” in *Proceedings of the 11th International Conference on Trust, Privacy & Security in Digital Business (TrustBus 2014)*, Munich, Germany.

Protsenko, M., and Müller, T. 2015. “Protecting Android Apps against Reverse Engineering by the Use of the Native Code,” in *Proceedings of the 12th International Conference on Trust, Privacy & Security in Digital Business (TrustBus 2015)*, Valencia, Spain.

Rakotondravony, N., Köstler, J., and Reiser, H. P. 2017. “Towards a Generic Architecture for Interactive Cost-aware Visualization of Monitoring Data in Distributed Systems,” in *Proceedings of the 17th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS 2017)*, Neuchatel, Switzerland.

Rakotondravony, N., and Reiser, H. P. 2016. “Visualizing and Controlling VMI-based Malware Analysis in IaaS Cloud,” in *Symposium on Reliable Distributed Systems (SRDS 2016), PhD Forum*.

Rakotondravony, N., Taubmann, B., Mandarawi, W., Weishäupl, E., Xu, P., Kolosnjaji, B., Protsenko, M., de Meer, H., and Reiser, H. P. 2017. “Classifying Malware Attacks in IaaS Cloud Environment,” *Proceedings of the Journal of Cloud Computing*.

Rauchecker, G., Yasasin, E., and Schryen, G. 2014. “A Decision Support System for IT Security Incident Management,” in *Proceedings of the 11th International Conference on Trust, Privacy and Security in Digital Business (TRUSTBUS 2014)*, pp. 36–47.

Reinfelder, L., Benenson, Z., and Gassmann, F. 2014. “Differences between Android and iPhone Users in Their Security and Privacy Awareness,” in *Proceedings of the 11th International Conference on Trust, Privacy & Security in Digital Business (TrustBus 2014)*, Munich, Germany.

Reinfelder, L., and Weishäupl, E. 2016. “A Literature Review on Smartphone Security in Organizations using a new theoretical Model-The Dynamic Security Success Model,” in *Proceedings of* *the 20th Pacific Asia Conference on Information Systems (PACIS 2016)*, Chiayi, Taiwan.

Reiser, H. P. 2017. “Towards Intrusion-resilient Security Monitoring in Multi-cloud Infrastructures,” in *Workshop on Security and Dependability of Multi-Domain Infrastructures, EuroSys 2017*, Belgrade, Serbia.

Richthammer, C., Kunz, M., Sänger, J., Hummer, M., and Pernul, G. 2015. “Dynamic Trust-based Recertifications in Identity and Access Management,” in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015)*, Vienna, Austria.

Richthammer, C., Netter, M., Riesner, M., Sänger, J., and Pernul, G. 2014. “Taxonomy of Social Network Data Types,” *EURASIP Journal on Information Security* (11).

Richthammer, C., and Pernul, G. 2016. “Explorative Analysis of Recommendations Through Interactive Visualization,” in *Proceedings of the 17th International Conference on Electronic Commerce and Web Technologies (EC-Web 2016)*, Porto, Portugal.

Richthammer, C., Weber, M., and Pernul, G. 2017. “Reputation-Enhanced Recommender Systems,” in *Proceedings of the 11th IFIP WG 11.11 International Conference on Trust Management (IFIPTM 2017)*, Gothenburg, Sweden.

Richthammer, H. 2014. “The Privacy Issues for Pseudonymised Customers in the Smart Grid,” in *Proceedings of the International Symposium on Engineering Secure Software and Systems (ESSoS 2014)*.

Richthammer, H., and Kesdogan, D. 2015. “Secure and Privacy Focused Customer Device Management in a Smart Household Environment,” in *Proceedings of the 30th IFIP International Information Security and Privacy Conference (IFIP SEC 2015)*.

Richthammer, H., and Reif, S. 2015. “Intrusion Detection in the Smart Grid based on an Analogue Technique,” in *International Workshop on Open Problems in Network Security (iNetSec 2015)*.

Russ, S., Reinfelder, L., Schankin, A., and Benenson, Z. 2017. “An Inquiry into Perception and Usage of Smartphone Permissions Models,” *(in progress)*.

Sänger, J., Hänsch, N., Glass, B., Benenson, Z., Landwirth, R., and Sasse, M. A. 2016. “Look Before You Leap: Improving the Users Ability to Detect Fraud in Electronic Marketplaces,” in *Proceedings of the 34th ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 2016).*

Sänger, J., and Pernul, G. 2014a. “Reusability for Trust and Reputation Systems,” in *Proceedings of the 8th IFIP WG 11.11 International Conference on Trust Management (IFIPTM 2014),* Singapore.

Sänger, J., and Pernul, G. 2014b. “Visualizing Transaction Context in Trust and Reputation Systems,” in *Proceedings of the 9th International Conference on Availability, Reliability and Security (ARES 2014),* Fribourg, Switzerland.

Sänger, J., and Pernul, G. 2017. “Interactive Reputation Systems,” *Proceedings of the Business & Information Systems Engineering*.

Sänger, J., and Pernul, G. 2016. “TRIVIA: Visualizing Reputation Profiles to Detect Malicious Sellers in Electronic Marketplaces,” *Journal of Trust Management* (3:5).

Sänger, J., Richthammer, C., Hassan, S., and Pernul, G. 2014. “Trust and Big Data: A Roadmap for Research,” in *Proceedings of the 1st Workshop on Security in Highly Connected IT Systems (SHCIS 2014),* Munich, Germany.

Sänger, J., Richthammer, C., Kremser, A., and Pernul, G. 2015. “Personalized Composition of Trustful Reputation Systems,” in *Proceedings of the 29th Annual IFIP WG 11.3 Working Conference on Data and Applications Security and Privacy (DBSec 2015),* Fairfax, VA.

Sänger, J., Richthammer, C., Kunz, M., Meier, S., and Pernul, G. 2015. “Visualizing Unfair Ratings in Online Reputation Systems,” in *Proceedings of the 23rd European Conference on Information Systems (ECIS 2015),* Münster, Germany.

Sänger, J., Richthammer, C., and Pernul, G. 2015. “Reusable Components for Online Reputation Systems,” *Journal of Trust Management* (2:5).

Sänger, J., Richthammer, C., Rösch, A., and Pernul, G. 2015. “Reusable Defense Components for Online Reputation Systems,” in

Proceedings of the 9th IFIP WG 11.11 International Conference on Trust Management (IFIPTM 2015), Hamburg, Germany.

Scheibel, M. 2014. "Die Analyse des Informationsflusses durch die Sensordaten und ihre langzeitliche Auswirkung auf den Datenschutz,", University of Regensburg.

Schryen, G., and Weishäupl, E. 2015. "IT-Sicherheit: Ökonomisch Planen und Bewerten," *Managementkompass* (2), pp. 17–18.

Sentanoe, S. 2017. "VMI based Honeypot (in progress),", University of Passau.

Sentanoe, S., Taubmann, B., and Reiser, H. P. 2017a. "Virtual Machine Introspection Based SSH Honeypot," in *EuroSys 2017, Poster*.

Sentanoe, S., Taubmann, B., and Reiser, H. P. 2017b. "Virtual Machine Introspection Based SSH Honeypot," in *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems (SHCIS 2017)*, Neuchatel, Switzerland.

Taubmann, B., Dusold, D., Frädrich, C., and Reiser, H. P. 2015. "Analysing Malware Attacks in the Cloud: A Use Case for the TLSInspector Toolkit," in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015)*, Vienna, Austria.

Taubmann, B., Frädrich, C., Dusold, D., and Reiser, H. P. 2016. "TLSkex: Harnessing virtual machine introspection for decrypting TLS communication," in *DFRWS EU 2016 Annual Conference*.

Taubmann, B., Huber, M., Heim, L., Sigl, G., and Reiser, H. P. 2015. "A Lightweight Framework for Cold Boot Based Forensics on Mobile Devices," in *Proceedings of the 10th International Conference on Availability, Reliability and Security (ARES 2015)*, Toulouse, France.

Taubmann, B., and Kolosnjaj, B. 2017. "Architecture for Resource-Aware VMI-based Cloud Malware Analysis," in *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems (SHCIS 2017)*, Neuchatel, Switzerland.

Taubmann, B., Rakotondravony, N., and Reiser, H. P. 2016. "CloudPhylactor: Harnessing Mandatory Access Control for Virtual Machine Introspection in Cloud Data Centers," in *The 15th*

IEEE International Conference on Trust, Security and Privacy in Computing and Communications (IEEE TrustCom-16).

Taubmann, B., and Reiser, H. P. 2016. "Secure Architecture for VMI-based Dynamic Malware Analysis in the Cloud," in *The IEEE/IFIP International Conference on DependableSystems and Networks (DSN 2016).*

Taubmann, B., Reiser, H. P., Kittel, T., Fischer, A., Mandarawi, W., and de Meer, H. 2015. "CloudIDEA: Cloud Intrusion Detection, Evidence preservation and Analysis," in *EuroSys 2015, Poster,* Bordeaux, France.

Vlad, M., and Reiser, H. P. 2014. "Towards a Flexible Virtualization based Architecture for Malware Detection and Analysis," in *Proceedings of the 1st Workshop on Security in Highly Connected IT Systems (SHCIS 2014),* Munich, Germany.

Webster, G. D., Kolosnjaji, B., Pentz, C. von, Kirsch, J., Hanif, Z. D., Zarras, A., and Eckert, C. 2017. "Finding the Needle: A Study of the PE32 Rich Header and Respective Malware Triage," in *Proceedings of the 14th Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA 2017)*, Bonn, Germany.

Weishäupl, E. 2017. "Towards a Multi-objective Optimization-model to Support Information Security Investment Decision-making," in *Proceedings of the 4th Workshop on Security in Highly Connected IT Systems (SHCIS 2017),* Neuchatel, Switzerland.

Weishäupl, E., Kunz, M., Yasasin, E., Wagner, G., Prester, J., Schryen, G., and Pernul, G. 2015. "Towards an Economic Approach to Identity and Access Management Systems Using Decision Theory," in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015),* Vienna, Austria.

Weishäupl, E., Yasasin, E., and Schryen, G. 2015a. "A Multi-Theoretical Literature Review on Information Security Investments using the Resource-Based View and the Organizational Learning Theory," in *Proceedings of the 36th International Conference of Information Systems (ICIS 2015),* Fort Worth, TX.

Weishäupl, E., Yasasin, E., and Schryen, G. 2015b. "IT Security Investments through the Lens of the Resource-based View: A new theoretical Model and Literature Review," in *Proceedings of the 23rd*

European Conference on Information Systems (ECIS 2015), Münster, Germany.

Weishäupl, E., Yasasin, E., and Schryen, G. 2017. “A Multiple Case Study on Information Security Investment Decision-making and Evaluation comparing the Consultant-Perspective with the Client-Perspective,” *Computers and Security*.

Werli, P., Popp, W., Frädrich, C., Sell, L., Fink, S. D., Bucher, A., and Marktscheffel, T. 2016. “Secure Smart Home,”.

Yasasin, E., Rauchecker, G., Prester, J., and Schryen, G. 2014. “A Fuzzy Security Investment Decision Support Model for Highly Distributed Systems,” in *Proceedings of the 1st Workshop on Security in Highly Connected IT Systems (SHCIS 2014),* Munich, Germany.

Yasasin, E., and Schryen, G. 2015. “Requirements for IT security Metrics - An Argumentation Theory Based Approach,” in *Proceedings of the 23rd European Conference on Information Systems (ECIS 2015),* Münster, Germany.

Yasasin, E., Weishäupl, E., and Schryen, G. 2017. “Review: Information Security Investments – A Multi-theoretical Foundation, Literature Synthesis and Research Issues,” *Submitted to the MIS Quarterly*.

Zach, J. 2014. “Design and Implementation of Integrated IaaS-Forensics for the Cloud,”, University of Passau.

Zach, J., and Reiser, H. P. 2015. “LiveCloudInspector: Towards Integrated IaaS Forensics in the Cloud,” in *Proceedings of the 15th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS 2015),* Grenoble, France.

Zhuang, Y., and Freiling, F. 2015. “Approximating Optimal Software Obfuscation for Android Applications,” in *Proceedings of the 2nd Workshop on Security in Highly Connected IT Systems (SHCIS 2015),* Vienna, Austria.

Zhuang, Y., Protsenko, M., Müller, T., and Freiling, F. 2014. “Measuring the Strength of Source Code Obfuscation Revisited,” in *Proceedings of the 1st Workshop on Security in highly connected IT systems (SHCIS 2014)*, Munich, Germany.

The Bavarian Research Alliance *FORSEC*, funded by the Bavarian State Ministry of Education, Science and Arts, is dedicated to research on security in highly connected IT systems. *FORSEC* is a joint research endeavour of ten research groups at four Bavarian universities (University of Regensburg, University of Passau, Technical University of Munich and Friedrich-Alexander-University of Erlangen-Nuremberg), has been coordinated by the University of Regensburg (Prof. Günther Pernul, Prof. Guido Schryen) and managed by Prof. Rolf Schillinger.

This book reports on the goals, projects and scientific achievements of *FORSEC*, which has published more than 100 research articles. It also demonstrates how *FORSEC* implemented three research clusters *PreSTA*, *STAR* and *CLOUD to* foster the collaboration between research projects and to achieve the overall research goals of *FORSEC*.