



SEcure Cloud computing for CRITICAL Infrastructure IT

Contract No 312758

Project Periodic Report Final Report 01/01/2013 – 31/12/2015

AIT Austrian Institute of Technology • ETRA Investigación y Desarrollo • Fraunhofer Institute for
Experimental Software Engineering IESE • Karlsruhe Institute of Technology • NEC Europe •
Lancaster University • Mirasys • Hellenic Telecommunications Organization OTE •
Ayuntamiento de Valencia • Amaris

Table of Contents

1	Executive summary	4
2	Summary description of the project context and the main objectives	6
3	Description of the main S & T results/foregrounds	10
3.1	WP2 – Requirements, Use cases and Legal	11
3.2	WP3 – Architecture, Specification and Design	15
3.3	WP4 – Cloud Operational Security and Design.....	20
3.4	WP5 – Cloud Analysis and Assurance	26
3.5	WP6 – Demonstration	29
3.6	WP1 – Project Management.....	33
4	Description of the potential impact, the main dissemination activities and the exploitation of results	37
4.1	Socio-economic impact and the wider societal implications.....	37
4.2	User and Advisory Board & Liaisons with other Projects	38
4.3	Standardisation activities	39
4.4	Dissemination Activities	40
4.5	Exploitation activities.....	40
4.5.1	Overview of exploitable products	42
4.5.2	Details of exploitable project results	44
5	Annex	54
5.1	CVs of Scientific Advisory Board Members	54
5.1.1	Prof. Dr.-Ing. Marcus Schöller	54
5.1.2	Dr. Markus Brunner	54
5.1.3	Professor Burkhard Schafer	54
5.2	Logos.....	55
5.2.1	Project Logo	55
5.2.2	Partner Logos.....	56

List of Tables

Table 1: Output Clusters	9
Table 2: Mapping of Architectural Framework Components to tasks	11
Table 3: SECCRIT Test Cases on First Iteration	30
Table 4: Demos, test cases, cloud platforms and RTD outputs	32
Table 5: Overview of exploitable products	43
Table 6: Exploitable Products	53

List of Figures

Figure 1: Project Objectives	8
Figure 2: SECCRIT Architectural Framework	10
Figure 3: Detailed Implementation of SECCRIT Demo 2	31
Figure 4: SECCRIT Logo Monochrome	55
Figure 5: SECCRIT LOGO, colour	55
Figure 6: AIT Austrian Institute of Technology GmbH	56
Figure 7: Amaris Technologies GmbH & Amaris Technologies Srl	56
Figure 8: ETRA Investigación y Desarrollo, S.A.	56
Figure 9: Fraunhofer IESE	56
Figure 10: Karlsruhe Institute of Technology	57
Figure 11: Mirasys Ltd.	57
Figure 12: NEC Europe Ltd., UK	57
Figure 13: Hellenic Telecommunications Organization S.A. (OTE)	57
Figure 14: Lancaster University	58
Figure 15: Ajuntament De Valencia	58

1 Executive summary

SECCRIT, SEcure Cloud computing for CRitical infrastructure IT, (contract number 312758) was a multidisciplinary research project bringing together ten partners from industry and research from Austria, Finland, Germany, Greece, Spain, Italy and the UK. SECCRIT's mission was to analyse and evaluate cloud-computing technologies specifically looking at security risks present in the context of critical infrastructures (CI). All activities were supported by a User and Advisory Board comprising stakeholders and practitioners alike. The project had a duration of 3 years (01/01/2013 to 31/12/2015) and an overall budget of about five million Euro. The consortium, coordinated by AIT the Austrian Institute of Technology comprised the following 10 Partners: AIT, ETRA Investigación y Desarrollo, S.A (ETRA, Spain), Fraunhofer IESE (IESE, Germany), Karlsruhe Institute of Technology (KIT, Germany), NEC Corporation (NEC, UK), Lancaster University (ULANC, UK), Mirasys Ltd. (MIRASYS, Finland), Hellenic Telecommunications Organisation S.A. (OTE, Greece), Ajuntament De Valencia (VLC, Spain) and Amaris (AMARIS, Austria).

Organised in seven work packages (incl. project management, and dissemination), the consortium defined five research objectives that represent the focus of work packages, but also research activities that required cross work package interaction. Individual results of various work packages have been presented together as **clustered outputs** to support exploitation and impact: **1. Techno-legal Guidance:** comprised a set of “stand alone” recommendations and guidance in legal issues and describes technical backgrounds, it is also integrated into various other outputs of other SECCRIT clusters. **2. Risk Assessment Methods and Tools:** comprised a vulnerability catalogue as input for a novel risk assessment methodology, developed in the project, and a tool in support of this methodology. Outputs have been contributed to ETSI Standardisation activities and evaluated as part of the SECCRIT demo activities by AIT, VLC and ETRA. **3. Policy Specification, Decision and Enforcement for Secure Data Handling in the Cloud:** comprised various add-ons for the usage control framework IND²UCE. IND²UCE components produced as part of SECCRIT activities have been made publicly available in the SECCRIT open source repository, all have been tested as part of the SECCRIT demonstrator evaluation. **4. Resilience Framework including Anomaly Detection in the Cloud:** included an anomaly detection framework (realised as Software prototype) and deployment functions for increased resilience of virtual resources. Results of this activity resulted in a number of high quality academic publications. **5. Tools for Audit Trails and Root Cause Analysis:** comprised software prototypes launched as “CloudInspector” (supporting generation of audit trails and root cause analysis) and “PoWerStore” (for secure storage of audit trails). Both components are being currently investigated for being commercially exploited for products and start-up companies, open source versions are available in the SECCRIT open source repository¹. **6. Cloud Assurance Profile Evaluation Method:** comprised the description of an assurance method together with corresponding proof of concept scripts. **7. Security Guidelines to support CI Providers in using the Cloud:** comprise a CI security requirements analysis and a “cloudification” guideline from CI IT – resulted from consolidating existing guidelines.

In total we have produced 31 deliverables. The seven output clusters have been evaluated in ten test cases. We have produced 38 peer-reviewed scientific papers of which 24 are collaborative. We have organised four User and Advisory Board workshops together with other EU initiatives and research projects, and we have organised a Dagstuhl seminar on “Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations”. Moreover, within the SECCRIT context, 11 student theses have been completed, and the SECCRIT project outputs

¹ <https://www.seccrit.eu/publications/source-code>

have contributed to several lectures. Partners have succeeded to launch five follow-up projects to continue work on the individual output clusters, and various commercial exploitations are planned and vivid interaction with the standardisation community has been established.

2 Summary description of the project context and the main objectives

Objective 1: *Establishment of legal guidance on SLA compliance, provision of evidence, and data protection for cloud services.* This objective was addressed in WP2. In the activities related to task T2.3 we established the legal fundamentals for technical SECCRIT project partners related to evidence law (relevant for SLA) and data protection issues. The outcome of this task has been reported at the end of M6 in deliverable D2.2 “Legal Fundamentals”. They contain practical examples, for people with a technological background rather than a legal one, for each of the identified legal issues. The results of our consultation with the national data protection authorities of Spain and Finland, where our demonstrators were located, can be found in D2.4. As a result from the review of year one, national data protection agencies have been consulted with particular regard to the demonstrator settings. A legal scholar was invited to become (together with two other senior scientific experts) a scientific advisor. He participated with the other two in the SECCRIT User and Advisory Board workshop and gave comments and recommendation. He also participated at a Dagstuhl seminar in April 2015, which was organised by SECCRIT partners, together with other legal scholars, to work on legal research questions in a multidisciplinary group. The legal team also provided the legal point of view in several deliverables and documented an overview of their work in D2.7 “Summary of Legal Aspects”. They have given an overview on ethics aspects in D2.8 “Final Ethics Report” and provided expert knowledge to D3.4 “Security Guideline”.

Objective 2: *Understand and manage risk associated with cloud environments.* The second objective has been addressed by WP3 in the first reporting period, but work on it continued in year two. The existing work was applied to a demo with partners from ETRA and VLC; the final results were documented in the corresponding reports on demonstration of SECCRIT RTD outputs. The work is now considered in ETSI GS NFV-REL 001². The vulnerability catalogue developed in WP3 (reported in D3.1 “Methodology for Risk Assessment and Management”) served as a basis for a number of other activities, such as a policy template catalogue and a policy elicitation method defined in Task 3.2. The method comprises seven process steps, including the identification of policy information sources (including our vulnerability catalogue from D3.1), a technology mapping, and the assessment of legal implications and technical feasibility. The work was documented in D3.3 “Policy Specification Tool” and in D4.4 “Policy Decision and Enforcement Tools”. The work on risk management was included in the demonstrators via corresponding evaluation activities and reported correspondingly. Furthermore, it was also incorporated into lectures and provided to students who were being trained in risk assessment of cloud applications. Our results have been documented in deliverable D3.2 “Policy Specification Methodology” and served as basis for our policy specification tool. In addition, the security policy catalogue supports the management of security demands and risks associated with cloud environments. Our vulnerability catalogue was furthermore considered in the work on deliverable D5.2 “Cloud assurance profile and evaluation method” which serves Objective 2 and 3.

Objective 3: *Understand cloud behaviour in the face of challenges.* This objective requires a thorough understanding of cloud environments and a common view on components and layers involved in a “cloudified” critical infrastructure IT environment. In WP5 an architectural framework – together with other WPs – was developed (reported in D5.1) within the context of an investigation on how reliability of audit trails can be achieved in a cloud environment. This work was being applied in all technical and techno-legal considerations and activities. An update was disseminated to stakeholders via the 6th SECCRIT newsletter and also documented in D5.3

² <http://bit.ly/1zxwAL0>

“Tools and Evaluation of Audit and Root-cause Tools”. Furthermore, in WP4 anomaly-detection challenges and a resilience framework for cloud environments was investigated and reported in deliverable D4.2 “Resilient Cloud Management”. This was supported by initial work on mitigation options via component deployment from WP5. WP5 contributed additionally to the topic of assurance evaluation methodology (D5.2). The methodology allows continuous aggregation of low level monitoring information and presenting it at assurance level. Via this various stakeholders can be informed on the impact of emerging challenging situations on assured security properties.

Objective 4: *Establish best practices for secure cloud service implementations.* This objective has been pursued intensely in Task 3.3. We have conducted a literature research to identify relevant resources targeting at cloud security and cloud security guidance. An analysis of how security requirements of critical infrastructure providers differ from security requirements of e.g. industrial stakeholders and how a potential difference can impact the taxonomy of a best practice guideline have been carried out. The analysis was supported with questionnaires which we distributed on our annual User and Advisory Board workshop and at other dissemination activities. The results of the questionnaire as well as a “cloudification” Security Guideline, which forms the output of the work on this objective, have been published in papers at conferences. These findings formed a core contribution to D3.4 “Security Guideline”. The practicality of this guideline was evaluated with ETRA and VLC within the “cloudification” of a traffic management system in the City of Valencia and documented in D3.4.

Objective 5: *Demonstration of SECCRIT research and development results in real-world application scenarios.* Two demonstration deployments have been undertaken in WP6: (i) using the cloud to support a traffic management system in the City of Valencia; and (ii) implementing a video surveillance system that monitors critical infrastructures with the support of cloud-based services. This objective has been addressed in deliverable D6.1 “Demonstrators Definition” in which we have not only detailed the demonstrator but clustered our outputs in artefacts usable by industry. This was derived from early exploitation activities (documented in D7.4). We have extended this when reviewing requirements and use cases in deliverable D2.6 “Update of requirements and use cases”. We defined test cases in which the individual outputs are being evaluated. The use cases, demos and output clusters have been presented at our User and Advisory Board workshop, which was organised in cooperation with other projects and organisations, to get some additional feedback from various stakeholders. This objective has been addressed also in this reporting period in deliverable D6.2 “Demonstrators validation” and in D6.3 “Report on validation results” in which we have reported on the outputs of our research clustered in artefacts usable by industry and validated in different test cases. For reality checks, the project outputs have been presented to our User and Advisory Board at two workshops³, which were organised in cooperation with other projects and organisations, to get some additional feedback from various stakeholders.

³ <http://thefutureofcloud.org/> and <http://www.info-com.gr/en/>




FIGURE 1: PROJECT OBJECTIVES

Figure 1 depicts the high level SECCRIT objectives and corresponding outputs and activities. They map to tasks and work packages. Overall, all of the SECCRIT objectives map to various research activities, tasks, and work packages. Individual outputs – e.g. policy specification, policy decision and enforcement together with the anomaly detection and technologies contribute to the cloud resilience management framework; hence we have defined the following output clusters:



Techno-legal Guidance: In line with high-level Objective 1 “Establishment of legal guidance on SLA compliance, provision of evidence, and data protection for cloud services”, this cluster comprises a set of “stand alone” documents for guidance in legal issues and describes technical backgrounds, and the cluster output was also integrated into various other outputs of other SECCRIT clusters.



Risk Assessment Methods and Tools: In line with Objective 2 “Understand and manage risk associated with cloud environments”, this cluster comprises a vulnerability catalogue as a plugin for a risk assessment tool and a methodology.



Policy Specification, Decision and Enforcement for Secure Data Handling in the Cloud: In line with Objective 2 “Understand and manage risk associated with cloud environments”, this cluster comprises various add-ons for the IND²UCE framework.



Resilience Framework including Anomaly Detection in the cloud: In line with Objective 3 “Understand cloud behaviour in the face of challenges”, the corresponding outputs include our anomaly detection framework including software to provision virtual resources in the system.



Tools for Audit Trails and Root Cause Analysis: In line with Objective 3 “Understand cloud behaviour in the face of challenges”, the main outputs for this cluster comprise software launched as “CloudInspector” (supporting generation of audit trails and root cause analysis) and “PoWerStore” (for secure storage of audit trails).



Cloud Assurance Profile Evaluation Method. In line with Objective 3 “Understand cloud behaviour in the face of challenges” the main outputs for this cluster comprise the description of an assurance method together with relevant proof of concept scripts.



Security Guidelines to support CI Providers in using the Cloud: In line with Objective 5 “Demonstration of SECCRIT research and development results in real-world application scenarios”, the outputs comprise rigorous stakeholder involvement, a method consolidation, and guideline documentation.

TABLE 1: OUTPUT CLUSTERS

All output clusters have been evaluated in line with Objective 5, and furthermore the work was documented in various scientific, peer reviewed publications.

3 Description of the main S & T results/foregrounds

In order to establish a common view amongst the consortium for identifying points in the cloud environment an architectural model was developed as a common framework. These points included where to place interfaces, monitoring points, and other functionalities which relate to outputs of individuals tasks. The proposed architectural framework was documented in D5.1 as well as in a scientific paper and aims at a more precise role distinction that allows for better security analysis, separation of responsibilities, identification of separate administrative interfaces, and for checking the influence and coverage of legal aspects. We hence start with explaining the architectural framework to give a structured overview of the individual activities.

Figure 2 shows the various abstraction levels of our architectural model, illustrates some components and represents a snapshot of ongoing work.




FIGURE 2: SECCRIT ARCHITECTURAL FRAMEWORK

In a workshop in September 2013 the individual activities in different technical work packages were mapped to the SECCRIT architectural framework.

	User Level	Service Level	Tenant Infrastructure Level	Cloud Infrastructure Level
T3.1 Methodology for cloud risk assessment and management	X	X	X	X
T3.2 Policy specification & mapping	X	X	X	
T3.3 Process-oriented security guideline	X	X		
T4.1 Anomaly based challenge onset detection			X	X
T4.2 Cloud resilience framework		X	X	X
T4.3 Policy decision and enforcement	X	X	X	X
T5.1 Cloud assurance evaluation		X	X	X
T5.2 Establishing audit trails			X	X
T5.3 Root-cause analysis in clouds		X	X	X

TABLE 2: MAPPING OF ARCHITECTURAL FRAMEWORK COMPONENTS TO TASKS

The matrix in Table 2 shows that the coverage of activities over the architectural framework levels and components is well-distributed amongst the technical work packages WP3, WP4, and WP5.

The project was characterised by a lot of collaborative scientific work which manifested itself in the publication track-record of the project, including engagements and organisation of scientific events e.g. seminar at Schloss Dagstuhl⁴. Another result of the collaborative work and good dissemination was the start of early exploitation activities. A fundamental pre-requirement of achievements was the establishment of a common view amongst the consortium for identifying points in the cloud environment where to place interfaces, monitoring points, and other functionalities which relate to outputs of individual tasks. Thus, an architectural model was developed as common framework in year one. This was published in various sources – see dissemination report. The efforts on creating a common view have been extended by defining output clusters and techno legal templates for the individual deliverables. This was in line with early exploitation activities and hence supports the real-world applicability and demonstrators of our work.

3.1 WP2 – Requirements, Use cases and Legal

The WP2 activities have contributed to milestone MS1, which consists of two main objectives: “Requirements for the entire project identified” and “Legal fundamentals defined”. The contribution to MS1 was obtained through the submission of two deliverables, D2.1 “Report on requirements and use cases” and D2.2 “Legal fundamentals”, both submitted at M6. Furthermore, three more deliverables have been submitted in M12, specifically D2.3 “Analysis of security-related aspects”, D2.4 “National Data Protection Consultation Results”, and D2.5 “Initial Ethics Report”. The DoW specifies a Quality Management Process and a Security Group which reviews

⁴ <http://www.dagstuhl.de/15151>

deliverables which could include e.g. personal data. The deliverables D2.3 - 2.5 support this activity.

Three tasks have been carried out for the achievement of the WP objectives: task T2.1 on the “Gathering requirements and user board”, task T2.2 concerning the “Specification of use case scenarios”, and task T2.3 about the “Establishment of legal fundamentals and provision of legal guidance”. There has been a fourth task planned in WP2, which will give legal support for our research activities. This task (T2.4) started after M6 and went on till the end of the project duration.

In particular, both the process of gathering and definition of the SECCRIT system requirements and the description of the two demo scenarios and the related use cases which should demonstrate the effectiveness of the project results have been carried out in tasks T2.1 and T2.2 and were presented at the first User and Advisory Board (UAB) workshop. There valuable feedback was collected.

Regarding the definition of the requirements in task T2.1, the outputs were classified into 11 clusters, and for each output a detailed set of several requirements were defined. The definition process used the Volere methodology and has involved all the consortium partners. Based on a web tool, the methodology specifies a three-stage process, consisting of requirements definition, requirement validation and requirements revision. The process has been iteratively repeated until a full agreement on the requirements was obtained. In total, 111 requirements have been identified.

Concerning task T2.2 on the definition of the scenarios and use cases, a modelling methodology has been applied. Again, an iterative process was adopted to define, to verify, and to agree upon scenarios and use cases. Deliverable D2.1 describes the two scenarios identified (i.e. “Storage and processing of sensitive video surveillance data” and “Hosting critical urban mobility services in the cloud”) and the related use cases.

In task T2.3, the goal was to identify the legal fundamentals that the whole project rests upon and to outline them in a way that is actually meaningful from the perspective of technological development. The activities focused on the investigation and the analysis of the main legal issues related to the adoption of cloud platforms for the implementation of critical infrastructures. Two aspects, in particular, have been faced: “evidence law” and “data protection law”. Generally speaking, the former focuses on the disclosing of the disputed measures and the proving of their authenticity and integrity during, for instance, a liability conflict or a dispute over SLA-compliance. The latter, in turn, is rather focused on minimizing the amount of data recognizable or known by others in order to protect the data subjects’ fundamental rights – a *conditio sine qua non* especially for moving critical infrastructure IT processes to the cloud. The respective fundamentals and their implications for technological design were reported in D2.2. Furthermore, a set of basic legal questions has been developed from the fundamental considerations for explicit examination within any technical deliverable. These basic questions cover both the evidence and the privacy/data protection dimension and allow for an initial estimation of the techno-legal aspects relevant to the respective technological artefacts. Due to the agreed-upon obligation to cover these questions within every technical deliverable and their inclusion in the deliverable template, they provide a reliable starting point for in-depth techno-legal considerations and discussions. Furthermore, and not to forget, this obligation also ensures that legal aspects were reliably taken into account during all activities of technological development, thereby fostering the SECCRIT idea of “regulatory compliance by design”.

Task 2.4 “Legal support for research activities” focused on legal modelling of nested three-party cloud settings and identification of legal implications; initiation of fundamental and paradigmatic (policymaking-oriented) considerations on the techno-legal dimension of Cloud Computing within critical infrastructure; participation in joint research activities with technical partners, particularly aimed at the preparation of interdisciplinary scientific articles and practice-oriented whitepapers on SECCRIT’s subjects; and more detailed investigations on aspects of legally relevant digital evidence, particularly including procedural aspects of multi-party settings of Cloud Computing. Nevertheless, several efforts have been spent in order to improve the SECCRIT requirements and define the use cases. In this respect deliverable D2.6 “Update of requirements and use cases” has been submitted.

In deliverable D2.3, activities related to the analysis of the security aspects concerning the SECCRIT solutions have been reported. Additionally to the originally, stand-alone purpose of this deliverable, together with D2.4 and D2.5, it supports the consortium’s Security Group (SG) in reviewing deliverables with higher levels of security requirements. The activity focused on the identification of potential security and privacy/data protection issues resulting from the implementation of demo scenarios. First, the security objectives have been identified by analysing the use cases defined for the two scenarios. In the second step, several security requirements have been defined for each of the assets involved in the demonstrators in order to assure that the security objectives have been thoroughly considered. Finally, a list of concrete security measures that should be taken into account has been identified.

Deliverable D2.4 describes the engagement process with the Spanish and Finnish Data Protection Authorities for their authorization of the activities related to the SECCRIT project. The interaction with the national authorities provided a set of recommendations that will be considered for the development of the solutions adopted in the demonstrators. There has been no “show-stopper” identified by the Spanish and Finnish Data Protection Authorities.

Finally, deliverable D2.5 summarizes the analysis of the ethical challenges that should be taken into consideration for all project activities. The analysis mainly provides a set of ethical safeguards to ensure that all project activities are in line with established ethical values.

At the very beginning of year two partners have worked on addressing the reviewer recommendations of National Legal aspects in D2.2 and D2.4

- KIT legal addressed this by compiling a questionnaire for the legal departments of all demo partners. Mirasys and ETRA, along with AMARIS, have already involved their legal departments.
- Result of this questionnaire were provided in D2.7 “Summary of legal aspects” in M36.

The purpose of the document was to give a more-in-depth analysis about the requirements gathered since the (early) beginning of the project in order to find out:

- If they are still relevant to the involved - and already identified - end users;
- If they are still properly aligned to project objectives.

Again, concerning to the use cases, an iterative process has been adopted to clearly defining the functionalities that are to be tested in the first iteration of demonstration activities and by also setting-up the different roles and responsibilities among partners of the SECCRIT consortium.

Techno-legal activities have contributed to D2.7 and D2.8 as well as to most technical deliverables as it ensured that legal and ethical aspects were carefully considered during

technology development. Furthermore, they contributed to the formulation of policy advice (also D2.7).

Results have been published in an economics/regulation-oriented paper (“An Agency Perspective to Cloud Computing”) accepted and presented to a conference (GECON2014); Talks at various dissemination events, including “BuildingTrustInCloud” and “BitsThatByte⁵”; several other papers were prepared with technical partners; including a techno-legal paper on the importance of transparency for data protection law and the respective problems in Cloud environments submitted to CLOSER.

The activity carried out in WP2 during the third year focused on preservation of evidence and data protection issues in cloud-based scenarios with specific regard to the critical infrastructure domain. Deliberations within the evidence law part showed that the legal proving position of the cloud user (as a plaintiff) is very weak. This is because it is too difficult for him to disprove the non-default of the cloud provider in cases where the provider, or a person whom the provider uses to perform his duty against the user, is at fault. Also in cases where the cloud user (as a defendant) himself is sued by the customer, due to a fault of the provider, the current legal evidence situation likely does not permit him to disprove a possible own default; this leads to the undesirable situation that the user loses the lawsuit. We therefore addressed this lack of proof of the cloud user by permitting him technically to access provider-independent proofs, which are neutral and truthful and help him to deduce a possible negligent action by the provider. Within the data protection law part, the cloud user was identified as the controller; thus, the user is the one who needs to fulfil the data subject's rights. As processing on behalf of the controller is generally a given in cloud computing scenarios, the cloud user remains responsible in cases where he is outsourcing personal data of the data subject. Thus, the user is still the one who has to fulfil the rights of the data subject. For the cloud user, however, this obligation is problematic because he cannot easily assess what the provider is doing with the data of the data subject and therefore can only trust the provider that he acts as agreed. Consequently, we again opted for technical solutions that provide the cloud user with real control and transparency in order to permit him an inside view of cloud internal management proceedings. Thus, the objectives set were fully achieved, as legal guidance for development activities in the project and legal fundamentals have been established. The accumulated knowledge on cloud-specific techno-legal challenges and respective solution approaches will be exploited as a basis for academic education, PhD theses, scientific articles, follow-up research projects, etc. Depending on the further development, provisioning of professional advisory opinions to relevant associations, regulatory bodies, and others might provide exploitation opportunities in the future.

The results of these activities were reported in deliverable D2.7 “Summary of legal aspects”, providing an in-depth analysis of evidence and data protection issues, critical infrastructure requirements, and recommendations. All technologies developed in SECCRIT were legally assessed, and possibilities to integrate them in the legal framework have been delineated. Thus, the objectives set were fully achieved, as legal guidance for development activities in the project and legal fundamentals have been established.

Furthermore, during the third year support for the research activities has been continued regarding the ethical issues that the SECCRIT outputs should meet. This support has produced deliverable D2.8 “Final Ethics Report”, providing an update of the original D2.5 “Initial Ethics Report”, laying out the ethical challenges of cloud computing, and creating a comprehensive documentation on the overall approach taken within the project in matters of ethical aspects.

⁵ www.bitsthatbyte.at

Additionally we have ensured via the role “Legal, Ethical, Privacy and Policy Issues Officer” during the whole project lifetime that all project activities do not infringe any ethical issues.

3.2 WP3 – Architecture, Specification and Design

Activities in WP3 have contributed to reaching milestone MS2 by finalizing work on deliverable D3.1 “Methodology for risk assessment and management”. Overall, three tasks have been planned in the course of this work package. Task T3.1 was completed in M12. In M6, we have started working on task T3.2, which deals with the specification of policies for cloud environments. This task finished in the beginning of the third year of the project duration (M30). Task T3.3, which addresses the specification of a process-oriented security guideline, started in the second year of the project (M20).

In task T3.1, a literature research has been conducted to analyse related work in the field of risk assessment. The major items are cloud-oriented risk analysis, cloud vulnerabilities, and cloud threats. The main results can be split up into two activities: 1. Understanding cloud-specific threats and vulnerabilities; 2. Developing cloud-oriented risk assessment processes and methods. As related work separates vulnerabilities and threats, we brought these two concerns together and proposed an overarching categorisation.

A cloud computing risk questionnaire was created, and all members of the User and Advisory Board were asked to participate in the survey. The goal of the questionnaire was to determine the perceived risks of organizations that either use or provide cloud computing. To tackle this, the questionnaire was split up in three parts. The first part consisted of questions regarding general information about the organization, views on risks and what kind of formal risk assessment within the organization exists. The second and third parts respectively asked questions about the assessment of short-term risks and organisation-evolution risks.

A vulnerability and threat catalogue has been documented. Input came from the literature research on cloud computing threats and vulnerabilities and from a structured analysis based on the SECCRIT architectural framework.

The catalogue is organized into several categories, such as the NIST essential cloud characteristics (e.g., on-demand self-service, broad network access, resource pooling), virtualization and organizational specific issues, physical cloud infrastructure, security and resilience control implementation challenges, and issues associated with contemporary cloud offerings. For each catalogue item, the primary security and dependability objectives it affects (confidentiality, integrity, availability) are highlighted.

Thereafter, a cloud-adopted risk assessment methodology has been devised, including a process that can be applied by organisations to determine the additional risk associated with using the cloud versus remaining with an existing non-cloud deployment. The SECCRIT cloud adoption risk assessment extends the Verinice-supported information security risk assessment process. We used a video surveillance scenario to analyse and illustrate how the approach could be implemented. The scenario was based on a video surveillance system that has several ICT assets with general potential to be migrated to the cloud (e.g., live video data, and anomaly detection services).

Several potential online risk metrics have been identified, which can be used as a starting point for developing a Risk Assessment as a Service (RAaaS) concept. Metrics provided by the cloud infrastructure to its tenants to improve risk management could be; for example, a notification of attacks detected by the cloud provider, the concrete mapping to specific physical resources for

redundant services, the current load of resources, or network anomalies. All results have been documented in deliverable D3.1 “Methodology for Risk Assessment and Management”.

Task T3.2, on policy specification and mapping, was looking at the different policy levels within the SECCRIT framework. On the one hand, these are Data Usage Control Policies (DUCP) that govern the access to data as well as the usage of these data after initial access has been granted. On the other hand, there are Cloud Resilience Policies (CRP) used for the protection of dynamic infrastructures against anomalies and threats. Within the work packages, we started with coordination activities that allow us to align the work on DUCP and CRP, and we also had a number of physical meetings focused on the topical work within these two areas.

A survey with focus on policy specification and context awareness has been conducted (in cooperation with WP4). It aimed at eliciting cloud security needs affecting the methodologies and tools for improving the security in cloud infrastructures within the SECCRIT project. The survey was offered to members of the User and Advisory Board (UAB), at the time comprising representatives from 46 companies working in the domain of critical infrastructure or that are cloud providers. In total, 60 persons from those companies were asked for participation. Nineteen participants started the survey and 15 of them completed the questionnaire. This survey revealed that there is a strong trend toward cloud computing; that is, many future services will be designed as cloud services. Two demands from the UAB relevant for task T3.2 were identified. First, the majority of the participants claimed security and privacy as well as loss of control over key IT systems and the infrastructure itself as their main concerns regarding service deployment in cloud infrastructures. Thus, enabling cloud users to specify and enforce their security demands in the form of security policies for protecting sensitive data and services is required. Second, although there is a need for end users to specify security policies, concerns were raised whether cloud users are capable of specifying correct security policies that do not jeopardize security. A user-friendly specification process that guides the end user is necessary.

A methodology for security policy specification was developed in task T3.2 to provide a basis for discussion within the consortium. Also a document for the specification of example policies was created. The example policies were elaborated according to the scenario descriptions of the two planned SECCRIT demonstrators. In addition to those exemplary security policies, Fraunhofer IESE and ULANC added potential security policies that their technologies developed in WP4 can enforce in excess of the demonstrator needs. To elaborate the policy document, a preliminary version of the policy specification methodology was applied. The example security policy document was sent to the demonstrator partners for further refinement and completion. A policy editor allows a stakeholder to specify a security policy. Different stakeholders shall be enabled to specify security policies, but they have different background knowledge about security, technology, or the specific application domain. Therefore, different stakeholder types need different levels of support and guidance during the policy specification process. As a first step toward user-friendly policy editors, usability concepts for improving the policy specification process were studied in a literature research. These usability concepts influenced the development of the policy specification tool in D3.3.

We started to analyse the transformation of high level security goals or security requirements into a machine readable format. As part of the SECCRIT architectural framework a policy-based resilience framework was developed that is expressive enough to cover DUCP and CRP. Within the latter policy domain, an existing framework covering network anomaly detection was extended to cloud environments. This includes the translation of cloud anomaly detection schemes into an Event-Condition-Action rule set. A set of policies and the associated mechanisms have been specified. Furthermore, we are investigating whether the concept of

Policy Patterns can be applied within cloud environments. It has been demonstrated that this concept allows us to flexibly combat different types of challenges by generalizing successful resilience solutions into reusable patterns of mechanisms. An offline analysis running in parallel to the operational cloud environment validates the patterns before they are deployed within the infrastructure.

Within task T3.2, we also considered Policy Refinement as a way to link the user space (i.e. DUCP) with the system space (i.e. CRP). The idea was to break high-level user policies further down and map them onto system or resilience policies (as appropriate). This approach is applicable to user policies that specify requirements on system protection, security, and resilience (in contrast to those solely dealing with data protection and security, which are directly implemented at the data layer). However, to ensure that there are no contradicting and conflicting policies, coordination is necessary. Our research into these aspects of task T3.2 is supported by Dr. Alberto Schaeffer-Filho, Associate Professor, Institute of Informatics, Federal University of Rio Grande do Sul (UFRGS), Porto Alegre, Brazil, who was a Visiting Researcher at ULANC.

The results were documented in deliverable D3.2 “Policy specification methodology” which was delivered in M18 and D3.3 “Policy specification tool” in M30.

Although task T3.1 finished already in M12 of the project, we extended our risk assessment to the Valencia traffic control system and published our results. Traffic control systems are usually running in dedicated data centres. The virtualization of traffic control components is expected to provide several benefits, including centralized service provisioning and management and simpler update procedures for the application. However, one key risk for a cloud-based deployment for traffic control systems is the potential impairment of availability. Our results have been published at the Reliable Networks Design and Modelling workshop (RNDM’14).

Based on the results of our cloud risk assessment, we also contributed to the challenge-fault catalogue in the ETSI Group Specifications on Network Function Virtualization⁶.

Our policy specification methodology has been addressed in task T3.2. To gain better insight into user needs and expectations regarding methodologies and tools for improving the security in cloud infrastructures, we have conducted a survey among the SECCRIT User and Advisory Board members. The survey provides a deeper understanding in cloud usage scenarios, such as private, public, hybrid, and community cloud deployment, their perceived advantages, and the users’ concerns about cloud adoption. Participants stated cost reduction, scalability, elasticity and especially availability, reliability, and resilience as the main potential advantages of cloud infrastructures. Security and privacy as well as loss of control over key IT systems and the infrastructure itself are main concerns stated by the participants of the survey. In the course of the survey, participants were also asked to rate specific cloud features, such as cloud-specific file systems or databases, mobile device integration, data usage control, and context-aware security. The survey also enquired security policies for the cloud. The survey results will guide our research in WP3 and WP4.

We developed a security policy template that contains several attributes for identification, classification, and explanation of security demands. To create a policy template catalogue, we defined a policy elicitation method that comprises seven process steps, including the identification of policy information sources, a technology mapping, and the assessment of legal implications and technical feasibility. The method was applied within the SECCRIT consortium and yielded 39 security policy templates in the categories cloud infrastructure (CI), tenant

⁶ <http://bit.ly/1zxwAL0>

infrastructure (TI), service provider (SP), and service user (SU), referring to the corresponding levels of our architectural model. In the technology mapping step, we mapped the security policy templates to technologies developed by the SECCRIT research partners. The policy specification and the policy template catalogue with the technology mapping show the interconnection and relation between all our technical work packages (WP3, 4, 5). Our results have been published on the Cloud Applications and Security Workshop (CAS'14).

Based on the policy elicitation method and the resulting security policy catalogue, we started the development of a user-friendly policy specification tool. The tool aims to provide a user-friendly way of security policy template instantiation. Thus, different users as well as administrators of the cloud infrastructure and its critical services shall be enabled to specify their security demands in form of security policies.

Most of the security policy templates we identified can be enforced by technologies developed within SECCRIT. We selected several security policy templates from our policy template catalogue that will be evaluated in our demonstrator test cases in WP6. Our industry partners identified the most relevant test cases for their demo scenarios, which are documented in deliverable D2.6.

The policy refinement, as part of the policy specification, has also been documented in task T3.2 as part of deliverable D3.2. The refinement involves decomposition, operationalization, deployment, and re-refinement respecting correctness, consistency, and minimality as refinement properties. The example scenarios “Availability of critical service” and “Resilience against DoS” have been illustrated. They directly relate to work carried out in WP4 to improve overall resilience management in cloud environments.

In task 3.2, we also focused on the design and development of the policy specification tool. Our goal was to develop a so-called PAP framework that supports the development of different policy editors tailored to the specific needs of the policy creator and the application domain. To this end, we first had to define a policy vocabulary model that combines security policy templates on a specification level with instantiation rules for producing machine-readable security policies that can be enforced by a policy enforcement framework (e.g., the IND²UCE framework). We used these security policy templates (documented in deliverable D3.2) that we elicited from the application domain of critical infrastructure services as input for our model.

Another important building block for the PAP framework is support for the different specification paradigms. One key challenge of policy editors is providing the appropriate support level for the policy creator. The spectrum of policy creators may range from a person trying to specify security policies the first time to an expert user who specifies such policies in his daily duties. Hence, depending on the knowledge level of the user and his task, we can provide different specification paradigms. A specification paradigm describes the way of interaction between the policy creator and the system in order to formulate the policy creator's security demand as a security policy in the PAP. In D3.3, we described the specification paradigms “Predefined Set of Security Policies”, “Selection from List of Predefined Policies”, “Security Policy Templates”, “Composition of Predefined Policy Blocks”, and “Plain Text Security Policy Specification”. The paradigms aim to align the degree of specification freedom to the knowledge level of the policy creator.

Another feature of the PAP framework is the support of different platforms, such as Android or Swing. The framework allows dynamic binding of the graphical interaction components and other platform-dependent features to the specification paradigms and the vocabulary model. The PAP framework can support multiple machine-readable security policy languages. Currently, only the IND²UCE policy language is implemented, but the framework supports other languages that

provide an XML Schema Definition (XSD). In summary, the PAP framework consists of four layers: platform, security policy specification paradigm, security policy vocabulary model, and target security policy language. We created two prototypical instantiations of the framework: A PAP on the Android operating system supporting the security policy templates specification paradigm and a PAP based on Swing for the composition of predefined blocks and the security policy templates specification paradigm. The PAP framework has been used to support the demos and test cases for the industrial use cases.

The general idea of the PAP framework was published at the 2nd International Conference on Information Systems Security and Privacy (ICISSP 2015). The policy specification tool (i.e. the PAP framework) and its approach are described in deliverable D3.3.

Regarding task T3.3, we started our work in project month 18. We conducted a literature research to identify resources targeting at cloud security and cloud security guidance. The German Federal Office for Security in Information Technology (BSI), for instance, publishes security recommendations and standards as baseline protection catalogues and modules. Valuable guidance for critical infrastructure cloud deployment, especially for public authorities, is also provided by the U.S. National Institute of Standards and Technology (NIST). Other important resources in the field of critical infrastructure protection and secure cloud computing are the European Union Agency for Network and Information Security (ENISA), the Cloud Security Alliance (CSA), and the Open Security Architecture (OSA).

In addition, we studied existing migration concepts and methodologies. We conducted two surveys among critical infrastructure providers, industry, and academic experts for highlighting and analysing differences between information security requirements of industry (non-critical infrastructure providers) and critical infrastructure providers. Overall, we collected 170 responses and got a clear statement that geolocation is of high importance for cloud computing in the critical infrastructure sector – which is in stark contrast to quite the opposite opinion from other industry sectors. We created a secure cloud migration taxonomy comprising 34 security controls used for evaluating existing guidelines. Our taxonomy answers three questions for the security controls: (1) Is the security control defined in the observed guideline? (2) Is the security control implemented as a process? (3) Does the control involve architecture or conceptual design?

We also realized that there is currently no security development life cycle that explicitly takes a “cloudification” scenario into account. Hence, we devised a novel approach, the “cloudification” Security Development Life cycle (CloudSDL). It aligns to a classical software development lifecycle and enriches security requirements by referencing cloud-related guidelines, best practices, and standards. We added legal support by providing a checklist for estimating data protection law requirements in cloud computing. The checklist is not exhaustive, but provides a first indication of which legal requirements need to be fulfilled.

Finally, there are clear links between the SECCRIT RTD outputs and the security guideline. For instance, the policy specification, decision and enforcement output cluster provides methods and tools to specify security demands in the form of security policies and technical measures to enforce them. The CloudSDL has been applied to the Urban Traffic Management System of the City of Valencia use case. In this use case, ETRA is the technology provider operating the virtual infrastructure and the Municipality of Valencia is the CI service user. To conclude, the CloudSDL provides support for secure migration of new or existing systems to the cloud (D3.4).

Besides the literature research, we prepared a questionnaire to support our activities to develop a security guideline for the migration of IT services to the cloud, with a focus on high assurance services for critical infrastructure IT. The questionnaire was offered to participants of the “Building

Trust in Cloud” event in Vienna, which was a joint event with our second User and Advisory Board (UAB) workshop as well as two other community events in Vienna. In addition, we distributed our questionnaire online to about 40 participants. Altogether we received about 110 responses, which sharpened our research in the field of security guidance. A position paper was prepared that included a security requirements analysis of the industry sector and critical infrastructure providers, based on the analysis of the questionnaire. The position paper also contained our approach for a combined information security guideline and legal guidance for CI providers.

WP3 contributed to all objectives in the SECCRIT project. Objective 1 “Definition of Legal Guidelines” is addressed by the legal checklist in the security guideline as well as the policy specification methodology, which includes the elicitation of legal requirements. Task 3.1 mainly addresses Objective 2 “Understand Cloud Computing Risks” by providing a methodology for cloud risk assessment and management. Objective 3 “Understand Cloud Behaviour” is supported by the developed tools. For instance, the policy specification tool supports the specification of machine-enforceable security policies that are finally enforced by the policy enforcement tools. Our security guideline and the risk assessment and management methodology have been evaluated in the Urban Traffic Management System use case. Hence, this directly relates to Objective 4 “Establish Best Practices for Secure Cloud Service Implementations” and Objective 5 “Demonstration of SECCRIT in real use cases”. Moreover, the policy specification tool supports both industrial demo scenarios.

3.3 WP4 – Cloud Operational Security and Design

Activities in this work package started by coordination and discussion at plenary meetings in Helsinki and Valencia. WP4 contributed to milestone MS2 and all three tasks have already started in the first project year. Deliverable D4.1 “Anomaly detection techniques” was submitted in M12 to the European Commission. This deliverable documents one part of the work that is to be done in task T4.1.

The first task on hand was to analyse the viability of the state-of-the-art anomaly detection techniques in elastic cloud deployment scenarios. The activities included but were not limited to a literature survey of non-cloud based anomaly detection techniques, setting up a test-bed, simulation of cloud traffic behaviour, injection of anomalies, and simulation of migration and observing how anomaly detection techniques are affected. This activity has contributed to overall objective 3, “Understanding cloud behaviour” and resulted in a methodology for an anomaly evaluation framework for cloud computing (refer to D4.1) which included reference implementations for state-of-the-art anomaly detection techniques. This framework will be extended to provide tools and mechanisms for anomaly detection in the cloud-operating context for D4.3. Further, to address the reviewers’ comments, the annotated dataset from our experimentation is continuously being used by the students from Lancaster University, as part of their Advanced Networking course.

The planning and setting up of a test bed was carried out in parallel by IESE and ULANC. In doing so, cloud solutions such as VMware, Xen (KVM) and OpenStack are analysed, and possible interfaces to retrieve information from the infrastructure, but also to interact with the infrastructure, are investigated. The KVM based test-bed at ULANC is extensively exploited to quantify the effect of elasticity on state-of-the-art anomaly detection techniques and to evaluate various components of our cloud resilience management framework for D4.2. The results are reported in deliverable D4.1 and D4.2.

For task T4.1, these activities included a literature survey of non-cloud based anomaly detection techniques, setting up a test-bed, simulating cloud traffic behaviour, injecting anomalies, simulating the migration, and analysing the viability of the state-of-the-art anomaly detection techniques in elastic cloud deployment scenarios. Based on these results, we argued that there is a need for robust, preferably real-time anomaly detection for cloud environments, where migration is a normal day-to-day operation. Therefore, we developed a technique for real-time anomaly detection based on the concept of data density. The density computation is expressed recursively, which makes the technique memory less, i.e., it does not need to store historical data. The lightweight nature of our approach makes it more suitable for deployment in cloud environments. Our results show that our proposed approach is effective in detecting high and low intensity network-level attacks with 98 percent accuracy.

This activity contributed to overall Objective 3 “Understanding cloud behaviour”, and resulted in a methodology for an anomaly evaluation framework for cloud computing (in D4.1) which includes reference implementations for state-of-the-art anomaly detection techniques. The framework was extended to provide tools and mechanisms for anomaly detection in the cloud-operating context for D4.3. In parallel, KIT reviewed related work on anomaly detection techniques, which contributed to deliverable D4.1. KIT also developed two tools (Distack and PktAnon), which are made available as part of our tool chain. These tools can aid further investigating the impact of anonymization on anomaly detection. The tools’ description was reported in D4.3. The annotated dataset from our experimentation is continuously being used by the students from Lancaster University as part of their Advanced Networking course. It has been made publicly available for the research community.

Moreover, we implemented our anomaly detector as a service to run in OpenStack using the Monasca API. Monasca is a scalable monitoring solution that leverages high-speed message queues and computational engines. The interaction of various components is reported in D4.3.

Building on deliverable D4.1 and on the collaboration between involved partners, research was carried out into a cloud resilience management framework, which was reported in D4.2. The first step was to understand the security and resilience requirements and to design a function that can translate security and resilience requirements into automatically deployable descriptions. In particular, a deployment function is needed that places instances of virtual resources according to the resilience and performance requirements of service users. This placement is done with consideration of the anomaly detection component of the framework, which ensures that any deviation from normal behaviour is detected. This component can further provide input for analysis to identify the root cause of anomalies in order to narrow the scope of remediation. The policy engine of our framework makes use of a policy-based decision to activate management actions on the cloud infrastructure, both on VM and host level as well as on functionality exposed by existing network elements.

In cooperation with WP5, we studied the deployment function to better understand its security and resilience requirements and to translate them into an actual instantiation on the cloud. Such cloud integration requires proper placement of instances and the implementation of resilience patterns in cloud deployment frameworks. In the context of the resilience framework, these requirements also include the supervision of deployed instances by anomaly detection. This work is complemented by extending the existing anomaly detection framework to provide triggers for policy enforcement for remediation controls against various challenges. The anomaly detection component of the framework allows plug and play of any suitable technique. We have chosen the one-class Support Vector Machine (SVM) algorithm for the implementation of our System Analysis Engine (SAE) as recommended in the previous review meeting. We have run various

experiments and finally implemented the tuned SVM for our SAE implementation. The results of our experiments, which were reported in D4.2, are very promising.

Work on the anomaly evaluation framework included creating a number of scripts and cloud configuration settings which were deployed within the testbed such as:

- Scripts for feature extraction, selection and aggregation of statistics from traces etc.
- Implementation of feature selection ranking from the system and network level
- Various attack scripts which are configurable to generate attack traffic with varying intensities
- Scripts to emulate migration of VMs across different nodes
- Reference implementation of state-of-the-art anomaly detection techniques.

In parallel to these activities, KIT collected and reviewed related work with respect to anomaly detection techniques, which contributed to deliverable D4.1. Currently KIT is investigating enhanced anomaly detection techniques and their suitability for running CI services in the cloud, using Distack and PktAnon as tools.

Furthermore, for task T4.2 (cloud resilience work), resilience patterns were sketched in the overall SECCRIT architectural framework in cooperation with WP5; this contains a deployment function that can understand the security and resilience requirements of the services to be deployed, and translate them into an instantiation on the cloud. This work was complemented by extending existing anomaly detection frameworks to provide triggers for policy enforcement for remediation controls against various challenges. This also includes the translation of cloud anomaly detection schemes into an ECA rule set.

Towards task T4.2 (cloud resilience framework), we presented a framework in D4.2, which models and then applies a resilience strategy (D^2R^2+DR)⁷ in a cloud operating context to diagnose anomalies. The framework uses an end-to-end feedback loop that allows remediation to be integrated with the existing cloud management systems.

The three main components of the framework are anomaly detection, deployment function and policy engine. The work on a deployment function was conducted in cooperation with WP5, with the goal of understanding the security and resilience requirements of the services to be deployed and translating them into an actual instantiation on the cloud. This requires proper placement of instances and the implementation of resilience patterns in cloud deployment frameworks. In the context of the resilience framework, these requirements also include the supervision of deployed instances by anomaly detection. This work is complemented by extending the existing anomaly detection framework to provide triggers for policy enforcement for remediation controls against various challenges. To sum up our resilience work, we report the following activities in this task: mapping of resilience strategy D^2R^2+DR onto the SECCRIT architecture; identification of logical interfaces, functions, and data flow of the Cloud Resilience Management Framework (CRMF) components; updating of the Data Collection Engine (DCE) to support various normalization techniques and integration of deployment functions with monitoring components of anomaly detection.

In task T4.3, Fraunhofer IESE set up a Cloud environment based on VMware. The testbed runs on two physical hosts managed by VMware vSphere and controlled by a VMware vCenter Server. The VMware vCenter Server provides a management interface for controlling virtual resources and their lifecycle. A Cloud storage infrastructure based on HBase and Hadoop was implemented

⁷ <https://wiki.ittc.ku.edu/resilinet/File:D2r2%2Bdr.png>

in the VMware environment, which can be used as a common Cloud application for test cases and also as a basis for integrating policy enforcement in the application.

The IND²UCE (Integrated Distributed Data Usage Control Enforcement) framework developed at Fraunhofer IESE was instantiated on and adapted to the VMware cloud environment. To this end, core components such as the decision engine (Policy Decision Point, PDP) were deployed to the Cloud environment. Moreover, components for interacting with the VMware cluster have been developed to realize policy enforcement. As the IND²UCE framework is based on the Event-Condition-Action (ECA) model, we converted system events from the cloud environment to IND²UCE-compliant events. This was done by a Policy Enforcement Point (PEP) that intercepts the cloud environment events through an interface of the VMware vCenter Server. We have chosen a generic approach and tested about 230 event types in our testbed. We assume that many more different system event types (overall about 700) are currently supported by our interception component, but not tested yet. A similar approach has been chosen for the IND²UCE actions that trigger system reactions in the Cloud environment. For this, an interface of the VMware vCenter Server is used. The actions are handled by a Policy Execution Point (PXP) that directly interacts with the management interface of the Cloud environment.

The policy engine uses the IND²UCE⁸ framework developed by Fraunhofer IESE to add overall resilience to cloud solutions by enforcing security policies at different levels of abstraction. The work carried out combined results from task T4.2, task T4.3 (WP4), and task T3.2 (WP3). Core components such as the decision engine (Policy Decision Point, PDP) have been deployed to the Cloud environment. Moreover, components for interacting with the VMware cluster have been developed to realize policy enforcement. To this end, a component to intercept relevant events in the system (Policy Enforcement Point, PEP), a component to perform actions in the system (Policy Execution Point, PXP), and components to retrieve additional information for the decision making (Policy Information Point, PIP), such as performance indicators or checking dedicated host criteria, have been developed. We also addressed the translation of cloud anomaly detection schemes into an ECA rule set, which is mainly researched in task T3.2 “Policy specification & mapping”.

Hence, we are able to retrieve information from events related to virtual machines (e.g., migration, lifecycle, powercycle), cluster (e.g., lifecycle, resources), physical hosts (e.g., host operations, networking), datastores, networking, as well as roles and permissions. Actions can be performed on virtual machines (e.g., power on/off, reset, reboot, relocate, clone, migrate), on the cluster as a whole (e.g., reconfigure) as well as on roles and permissions (e.g., set/reset/remove entity permission). All these events and actions can be described within the event or action part of our policy language and finally used to specify security demands.

For D4.1, the activities included but were not limited to a literature survey of non-cloud based anomaly detection techniques, setting up a test-bed, simulate cloud traffic behaviour, inject anomalies, simulate the migration and observe how anomaly detection techniques are affected. The results of this work are published in IEEE CloudNet 2014.

Building on the work from D4.1, and on the collaboration between involved partners, research is carried out into a resilience management framework, which is reported in D4.2. The first step was to understand the security and resilience requirements and to design a function, which can translate security and resilience requirements into automatically deployable descriptions. In particular, a deployment function is needed that places instances of virtual resources according to

⁸ The IND²UCE framework received the Innovation Prize of the European Association of Research and Technology Organisation EARTO (October 2014).

the resilience and performance requirements of service users. This placement was done with consideration of the anomaly detection component of the framework that ensures that any deviation from normal behaviour is detected. It can further provide input for analysis to identify the root cause of anomalies in order to narrow the scope of remediation. Currently, we are investigating different anomaly identification/classification techniques such as “Fuzzy rule based (FRB)” to empower (Fine Grain Analysis) engine of our framework. Moreover, the remediation actions are triggered using a policy engine. The policy engine makes use of a policy based decision, to activate management actions on the cloud infrastructure both on the VM and host levels as well as functionality exposed by existing network elements. The work was submitted to “e&i Elektrotechnik und Informationstechnik” Springer Journal.

Furthermore, in parallel KIT collected and reviewed work with respect to anomaly detection techniques, which contributed to deliverable D4.1. KIT investigated enhanced anomaly detection techniques and their suitability for running CI services in the cloud, and the use of Distack and PktAnon⁹ as tools which was reported in D4.3

The policy decision and enforcement task, task T4.3, started to analyse possibilities about where to place monitoring and enforcement components best in cloud environments and what kind of adaptations are needed. This work was carried out by Fraunhofer IESE. Within the task, virtualisation solutions, such as VMware and OpenStack, and their available interfaces to retrieve information from the infrastructure were analysed. We also investigated how to interact with the infrastructure to enforce security and resilience policies. To this end, the enforcement framework IND²UCE has been extended and adapted for cloud solutions based on VMware. Furthermore, the framework has been applied to a cloud storage infrastructure based on HBase and Hadoop. We directly integrated our enforcement components by changing the source code of HBase, contrary to the VMware enforcement, where we use the provided interfaces. We can deploy data usage control policies and apply our enforcement to data in transit. For example, we can modify data in transit to anonymize data and preserve privacy. We performed several measurements to evaluate the overall impact of our policy enforcement framework to such a performance critical cloud storage infrastructure. We started to interconnect the enforcement in the cloud infrastructure (VMware) and the enforcement on the cloud storage.

Task 4.3 combined results from task T4.2 and task T3.2 (WP3). Core components, such as the decision engine (Policy Decision Point, PDP) and a management component (Policy Management Point, PMP), have been deployed to the cloud environment. Moreover, components for interacting with the VMware cluster have been developed to realize policy enforcement. To this end, a component to intercept relevant events in the system (Policy Enforcement Point, PEP), a component to perform actions in the system (Policy Execution Point, PXP) and components to retrieve additional information for the decision making (Policy Information Point, PIP), such as performance indicators or checking dedicated host criteria, have been developed. We also addressed the translation of cloud anomaly detection schemes into an ECA rule set, which is mainly researched in Task T3.2 “Policy specification & mapping”. Task T4.3 also comprised the development of enforcement components for a cloud storage infrastructure based on HBase and Hadoop, which runs in the VMware cloud environment. We directly integrated our enforcement components by changing the source code of HBase, contrary to the VMware enforcement, where we use the available interfaces. With our components in place, we can deploy data usage control policies and apply our enforcement to data in transit. For example, we can anonymize transmitted data to preserve privacy.

⁹ <http://www.tm.uka.de/software/pktanon/>

The proof of concept of these technologies was implemented for the SECCRIT demo test cases in WP 4. These test cases are validated and evaluated in D6.2 and D6.3. There have been some further extensions to the IND²UCE framework for better supporting and illustrating the demo test cases. For instance, we developed two policy execution points (PXPs) to directly interact with the UIs from ETRA and MIRASYS. Hence, the behaviour of the IND²UCE framework can be visualized in their graphical interfaces such as the ETRA I+D Alert Monitor.

The policy decision and enforcement using IND²UCE framework was implemented within the VMWare cluster under the specific test case "Growth in resource consumption-TC004". In this test case, the CPU load of a virtual machine is monitored and on high load, an additional CPU is allocated to the VM. This test case corresponds directly to Story 3 "Data not available due to a malfunction or misbehaviour" of UC-002 in D2.1. IND²UCE increases the resilience of the critical service by managing resources based on the specified security policies. A second application is test case "Dedicated Host", in which the IND²UCE framework prevents the operation of two critical services on the same physical host. The test case corresponds to Story 2 "The misbehaving politician" of UC-001 in D2.1, where sensitive data leaked from the system. The physical separation of the services should prevent such a data leakage.

The Deployment Function as part of the Resilience Framework (described in D4.2) was implemented for the specific test case "Failure Recovery" (TC007) of the "Storage and Processing of Sensitive CCTV Data" use case defined in SECCRIT. To this end, the support of an active-backup resilience pattern was integrated into the OpenStack orchestration workflow, and it was shown to work with the virtual servers of the project demo partner Mirasys, adapted to support failover. The concept of the Deployment Function addresses mainly Objective 2 'Understand and manage risk associated with cloud environments', as it is a tool to manage resilience and - in a certain sense - to define policies (in an abstract fashion, using, e.g., target availability values). As part of the Resilience Framework, and potentially being used to place specific parts of a deployment under the vigilance of Anomaly Detection, it also partially contributes to Objective 3 'Understand cloud behaviour in the face of challenges'. The Deployment Function was demonstrated in test case 007 defined in D2.6. This demonstration has been documented in D6.2 and D6.3. It has been exploited scientifically in terms of publications, as well as providing the basis for contributing to current standardization efforts in ETSI NFV, thus increasing the profile of NEC in this business-relevant standardization body. The demonstrator developed in the SECCRIT context will be used to show the concept to interested business units, and it is considered to be combined with other results in the NFV context, such as the open source initiative OPNFV.

The main achievements in the third year was the development of an Evaluation Framework for Anomaly Detection Techniques for Cloud Computing Experiments (EFADT-C2X), Cloud Resilience Management (CRM), and Anomaly Detection-as-a-Service (ADaaS), and the provision of Policy Decision and Enforcement components, which have been included in the deliverables D4.1, D4.2, D4.3 and D4.4, respectively. The anomaly detection framework was implemented for specific test cases; "Database that grows unexpectedly" (TC005) and "Network pattern changes" (TC006). For TC005, our framework detected an unexpected growth in the DB servers as anomaly and generated an alert to minimise the impact on the overall system. Similarly, for TC006 an anomalous pattern "network traffic between communications server and devices on public areas" is detected, and an alert is generated for an operator to invoke further remediation. Both of these test cases relate to Demo 2 and correspond to Story 3 "Data not available due to malfunction or misbehaviour" of UC-002 in D2.1.

The results of our demonstrations are documented in D6.2 and D6.3.

3.4 WP5 – Cloud Analysis and Assurance

WP 5 focused on cloud assurance and analysis. All WP5 partners actively contributed to the requirements engineering task within WP2 at the beginning of the project. Task T5.1, which started in M6, works on cloud assurance and has defined its scope in the form of a pending position paper. This work was started with some delay due to initial staffing problems. The task's goal was to find how to derive a uniform assurance level for a cloud service (e.g. if considering common criteria) if individual underlying components have different assurance levels and given that these components change over time. Therefore, existing assurance approaches are reviewed and analysed for how well they succeed and how well a common criteria approach fits into all of this. As part of the task, a liaison and researcher exchange was set up with colleagues from the CUMULUS project.

The tasks on auditing (T5.2) and analysis tools (T5.3) also started in the first year, and they worked jointly to define a common architectural framework for the overall project. It was recognised that only with this framework the tasks could start designing purposeful tools for their respective activities. This architectural framework was described in detail in deliverable D5.1 "Audit Trails".

Measuring and auditing of all systems involved in executing critical infrastructure services is essential to establish trust of the service operator into the infrastructure. Also, as laid out in D2.2 on legal fundamentals, we need to allow for legally compliant cloud usage and legally enforceable cloud contracts/SLAs. In order to accomplish this task, the SECCRIT project's architectural framework clearly illustrates the various stakeholders, actors, roles, and responsibilities as well as suitable abstractions to separate these entities. In addition, multiple views onto this architectural framework were developed, each of which focuses on a particular aspect of cloud based critical infrastructure services:

1. A multi-provider and multi-tenancy view: a particular security and dependability challenge stems from the operational mode of public cloud offerings. Services of multiple tenants are hosted co-located on a shared infrastructure to realize cost savings. This concept is called multi-tenancy of clouds. On the other hand, a service provider can request virtual resources from a set of cloud infrastructure providers to run its service. This concept is called a multi-cloud deployment. Such deployments pose their own security and dependability risks.
2. A network access view: the way information sources and service consumers connect to the service instance changes significantly compared to the traditional service model. Instead of a dedicated infrastructure, these entities connect via a public wide area network (the Internet) to the service instances. Moreover, these service instances are not fixed to a particular network location but can be migrated within or even across data centre boundaries.
3. A management view: the day-to-day operational management of the critical infrastructure services deployed on the cloud is crucial to their security and fault management. The management system controls the complete lifecycle of the service from its onset, instantiation, and execution, to its termination. In particular, we have focused on two specific aspects of management:
 - Monitoring: supervision of the service components with respect to performance, availability, and reliability.

- **Auditing:** generation of logging information to trace back the cause of a degradation or even failure of service provisioning. This auditing needs to respect administrative domains while providing a high degree of trustworthiness to all stakeholders of the system.

The architectural framework was published and presented at the IEEE CloudCom conference in December 2013. Then, as a first result of these views on the architectural framework, we conducted a risk assessment in task T3.1, detailing new threats and vulnerabilities introduced by moving critical infrastructure services to the cloud. And as previously indicated, this architectural framework was defined as the foundation for all technical tasks of the SECCRIT project. The various contributions from SECCRIT partners was aligned to this framework and we expect significant simplifications in designing the use case demonstrators by adhering to it.

For the second year, within WP5, we developed suitable tools for auditing and fault identification and to identify/develop suitable cloud assurance methods. Based on the management view of the architectural framework, we have already identified the interfaces and parameters which are essential to log in order to analyse the system fault in case of a service failure. We based our analysis on OpenStack, an open source cloud platform, which we can modify to fit our needs. In addition, we have already developed and evaluated a distributed storage mechanism which provides increased reliability and trustworthiness. It has been developed for secure, robust, and privacy-preserving storage of information. This tool can be used as a storage solution for sensitive service data as well as for an external trusted auditor. This auditor is envisioned to store audit information of the individual levels of the architectural framework (horizontal auditing) or across several levels for individual services (vertical auditing). The information stored by this external auditor is designed as a write-only tool during normal operation which guarantees that data is never altered, manipulated or deleted. Information from this storage is only read in case of a service failure and a dispute on liability, and it provides “digital evidence”. This is, as outlined in WP2, an important precondition for enforceable cloud contracts and is thus indispensable for broader adoption of cloud computing, especially in the critical infrastructure domain. Mechanisms to increase the trustworthiness of the information sent to this auditor is a further central aspect of our upcoming activities.

In task T5.1 (Cloud Assurance Evaluation), a framework for assurance evaluation and monitoring was developed in order to better understand and manage the risks associated with cloud environments. The work includes a survey of related work and an initial conceptual plan, published at IEEE FiCloud. In a second step we have designed the framework in more detail considering state-of-the-art monitoring tools for acquiring relevant input data sets across all multiple distinct layers as identified in WP3. Based on this overview, and on input from WP6 partners, a first catalogue of measurable assurance properties in a cloud context was developed which also served as the basis for future evaluations of the framework in the context of the test cases defined in D2.6.

Furthermore, assurance dependency policies supported by these properties were investigated and an algorithm for performing a complex system, service and information assurance analysis developed and published at the 2014 IEEE International Conference on Cloud Computing Technology and Science (IEEE CloudCom). This tool aims to increase awareness and trustworthiness of security and privacy aspects for cloud environments by providing a consolidated view of the assurance level of a given cloud infrastructure or a given application deployed in a cloud system. An evaluation based on feedback from the demonstration partners showed the applicability of the framework to relevant use-cases. The results of this work and of task T5.1 in general were documented and summarized in D5.2 in M24.

In the context of the work described above, the collaboration with the CUMULUS project was continued. A researcher participating in CUMULUS, Maria Krotsiani, visited AIT for four months, leading to a mutually beneficial exchange of research ideas and expertise. This addressed the reviewer comments twofold: first by considering existing certification schemes as used in CUMULUS and second by collaborating with other relevant projects.

Tasks T5.2 and T5.3 had the goal of developing tools for the establishment and analysis of audit trails, in order to better understand cloud behaviour in the presence of challenges. After having created the common framework necessary to guide the design of the individual tools (e.g., their interfaces and interactions) and the further development of these tools, the conceptual work was finished in 2015 with the tools reaching a state ready for demonstration and integration.

The architectural framework developed in SECCRIT allows four different views (e.g., physical view, management view, or security view). With respect to the auditing view, analysing the framework and its different layers (user, service, tenant, and cloud infrastructure) led to the conclusion that increased transparency is needed specifically for the cloud infrastructure level in order to allow for trustworthy auditing by higher layers. In addition, any viable auditing toolset needs appropriate logging and reliable and secure storage features in order to guarantee the integrity of the gathered audit data.

Based on these conclusions, several components have been developed catering to these needs. One of these components is an instantiation of an Independent Transparency-as-a-Service-Framework (named “CloudInspector”) that provides a separate view on the infrastructure independent from the information provided by the cloud infrastructure management or provider. The framework is documented in an update to D5.1. It is based on information gathering modules installed on the physical hosts of the infrastructure, that is, in parallel to the virtualized nodes. Thus, these modules provide a view from outside the virtualized environment, which is collected by an independent system and can thus be queried by tenants or higher layers without having to rely on information provided (and possibly filtered) by the Cloud Infrastructure Management System (e.g., OpenStack or VSphere).

However, it is not generally in the interest of cloud infrastructure providers to expose their configuration and operation details to public view. To resolve this issue (need for auditability vs. preservation of confidentiality), a commitment scheme has been developed and implemented that allows for the confidential and non-reputable storage of auditing information. This information is committed in a form that is not directly accessible but can be revealed when the need for an audit arises. Moreover, it cannot be changed after the original event has been committed.

With respect to the storage of audit trail data, a trustworthy and highly reliable solution for this component is necessary. To address this requirement, different storage principles have been developed, with a particular focus on PoWerStore. This protocol developed in SECCRIT is a Byzantine Fault-tolerant (BFT) storage solution that - compared to existing protocols - is highly efficient (i.e., it achieves optimal latency) for the features it offers, namely strong consistency and high availability. It achieves this degree of efficiency by using Proofs of Writing (PoW), which enable writing back metadata instead of the data itself (writing back data is a cornerstone for robustness and fault tolerance). An implementation of this protocol (“Hybris”) was coupled with the aforementioned CloudInspector using a standardized interface protocol.

Thus, the output of WP5 has contributed to three of the objectives set for the SECCRIT project: The output cluster “Assurance” has contributed mainly to Objective 3 “Understand cloud behaviour in the face of challenges”, as it provides a methodology and tools for continuous assurance evaluation. The output cluster “Tools for Audit Trails” has also contributed mainly to

Objective 3 “Understand cloud behaviour in the face of challenges” by providing tools to transparently monitor and log key characteristics and the operational state of a cloud deployment. In addition, the Tools for Audit Trails can serve to mitigate legal risks by providing independent and reliable audit trails, and it can also be used to monitor the compliance with set policies, thus also contributing to Objective 2. Finally, both output clusters have contributed to Objective 4 “Establish best practices for secure cloud service implementations”, since the provisioning of the developed methodologies and frameworks is recommended by SECCRIT as a best practice, and the key features checked by the Assurance Framework and the Tools for Audit Trails can be used to monitor the implementation and following of best practice guidelines, for example, by observing the locality of VMs.

The results of WP5 were demonstrated in the test cases (TC-nbr) defined in deliverable D2.6, namely in the TC-002, TC-005, TC-007, and TC009. The demonstrators have been documented in D6.2 and D6.3.

While all results have already been exploited scientifically in terms of publications and further research, the TAT component CloudInspector is currently also being investigated as a potential foundation for a spin-off of the Karlsruhe Institute of Technology (KIT), in addition to plans of further use or development in other research projects. The storage component Hybris and its underlying protocols will be further investigated in other EU projects, such as the H2020 project TREDISEC, and the knowledge gained has been transferred NEC-internally to interested business units, with the goal of exploiting the results in the form of NEC cloud solution products. The Assurance Framework will be developed further in future research projects, such as CREDENTIAL.

3.5 WP6 – Demonstration

Task T6.1 started with the work package kick-off in the plenary meeting held in Valencia in M11. The involved partners agreed on the work ahead in order to start producing D6.1 “Demonstrators definition”, as the task leader, started the preparation of the report. Furthermore, the partners agreed to have a catalogue of threats to be used in the demonstrator tests, which will complement the use cases and requirements defined at the beginning of the project, and most probably will be a basis to produce a scientific paper.

Four activities have been running during the lifetime of the SECCRIT project. Tasks T6.1 was dealing with the detail specification of the demonstrators, T6.2 where lead technological partners has supported end users in the definition and implementation of the demonstrators, T6.3 which covers all implementation issues of Finnish demo and T6.4 which covers Valencia implementations, and finally task T6.5 dealing with the analysis of the results and conclusions.

At the initial phase of WP6 it was agreed to have two iterations to validate technologies, a detailed plan was created for that purpose.

One of the main challenges in this period was to derive realistic testing scenarios, in that sense the requirements and use cases defined by the beginning of the project were taken into account together with the RTD Outputs.

In terms of implementation work, two different environments have been defined to test and validate the technologies: one related to VMWare technology supported by QloudWise, the commercial cloud of AMARIS; and another one based on OpenStack. The VMWare environment was quickly set up and available was ready for use for the demonstrations, whilst the setting up of

the OpenStack environment had taken more time than expected, mainly due to unforeseen issues while deploying.

Table 3 summarises the test cases taking place in the first iteration:

Test Case (TC) ID (identifier)	Description	Expected reaction	Platform	RTD (Research and Technological Development) Outputs used
Demo 1: Storage and processing of sensitive data				
TC-002	Dedicated host, i.e. anti-affinity of virtual machines	Migration VM (Virtual Machine)	VMWare ¹⁰	- Policy Specification, Decision and Enforcement - Tools for Audit Trails and Root Cause Analysis
TC-007	Failure recovery of a virtual machine with minimum interruption to a service	VM Replacement	OpenStack ¹¹	- Tools for Audit Trails and Root Cause Analysis - Resilience Framework with focus on Deployment Function - Assurance Framework
Demo 2: Hosting critical urban mobility services				
TC-001	Risk assessment of mobility services in the cloud	Operators / Owners Awareness	Independent	Risk Assessment
TC-003	Lost network connectivity for the database VM	Notification	VMWare	Policy Specification, Decision and Enforcement
TC-004	Unexpected growth in resource consumption on host	Notification and dyn. resource adaptation	VMWare	Policy Specification, Decision and Enforcement
TC-005	Database grows unexpectedly	Notification	OpenStack	- Resilience Framework with focus on Anomaly Detection - Assurance Framework
TC-006	Network pattern changes	Notification	OpenStack	Resilience Framework with focus on Anomaly Detection

TABLE 3: SECCRIT TEST CASES ON FIRST ITERATION

Furthermore a set of Virtual Machines (VMs) have been set up and configured to emulate the behaviour of both Finnish and Spanish demonstrators, they emulate mission critical services, in the first case storage of sensitive video data, and the second related to urban mobility services. The VMs are already uploaded and running in the VMWare environment and integrated the process to integrate the RTD Outputs, and make the first tests. Figure 3 shows a high level view of the VMs deployed in Demo 2, where three VMs emulate the data processing and storage of the Traffic Control Centre of the City of Valencia, and (emulated) sensor data is constantly uploaded to the cloud.

¹⁰ VMware virtualizes computing, from the data center to the cloud to mobile devices, to help our customers be more agile, responsive, and profitable. More information can be found at: <http://www.vmware.com/>

¹¹ OpenStack is a free and open-source cloud computing software platform. Users primarily deploy it as an infrastructure as a service (IaaS) solution. More related information can be found at: <http://www.openstack.org/>.




FIGURE 3: DETAILED IMPLEMENTATION OF SECCRIT DEMO 2

Even if it broads the initial scope of this WP2 deliverable we reported WP6 activities there since there are no more official deliverables in WP6 until the end of the project.

Hence, testing and validation of RTD outputs has been done in two iterations. While seven test cases (TC-nbr) in Iteration 1 were already defined in D2.6 (TC-001 – TC-007), later on a total of three new test cases were added for Iteration 2: TC008, TC009, and TC010. Further, TC-002 was refined during test Iteration 2.

Table 4 shows all of the test cases that were run and evaluated:

Test Case ID	Description	Expected reaction	Platform	RTD Outputs used	
Demo 1: Storage and processing of sensitive data					
TC-002	Dedicate host, i.e., anti-affinity of virtual machines	Migrate VM and provide an independent view of current situation	VMware		Policy Specification, Decision and Enforcement
					Tools for Audit Trails and Root Cause Analysis
TC-007	Failure recovery of a virtual machine with minimum interruption to a service	VM Replacement	Openstack		Tools for Audit Trails and Root Cause Analysis
					Resilience Framework with focus on Deployment Function
					Assurance Framework
TC-008	Asserting Right of access (Data Protection Law) — Geo location of personal data	Real-time inside View	Openstack		Tools for Audit Trails and Root Cause Analysis
					Legal Guidance
Demo 2: Hosting critical urban mobility services					
TC-001	Risk assessment of mobility services in the cloud	Operators/Owners Awareness	Independent		Risk Assessment
TC-003	Lost network connectivity for the database VM	Notification	VMware		Policy Specification, Decision and Enforcement
TC-004	Unexpected growth in resource consumption on host	Notification and dynamic resource adaption	VMWare		Policy Specification, Decision and Enforcement
TC-005	Database grows unexpectedly	Notification	Openstack		Resilience Framework focus on Anomaly Detection
TC-006	Network pattern changes	Notification	Openstack		Resilience Framework with focus on Anomaly Detection
TC-009	Legal evidence provision for proving negligent behaviour	Audit Trails	VMware		Tools for Audit Trails and Root Cause Analysis
					Legal Guidance
TC-010	Real-time monitoring of issues in the cloud	Notification via GUI	VMware		Policy Specification, Decision and Enforcement
					Resilience framework

TABLE 4: DEMOS, TEST CASES, CLOUD PLATFORMS AND RTD OUTPUTS

All this work has been reported in two different deliverables, both delivered on time:

- D6.1 Demonstrator definition
- D6.2 Demonstrators validation
- D6.3 Report on validation results

3.6 WP1 – Project Management

WP1 activities started with organising the first plenary meeting (kick-off meeting) in January 2013 in Vienna. This lasted three days, and we discussed the *modus operandi* (e.g. review process, reporting) of the project (and how to document this in the handbook), how to set up the website, the project portal, and we decided on a logo (see 5.2.1). At the end of M3 D1.1 “Project Handbook and Website” was delivered to the European Commission. Similar efforts were also put into organising the other three plenary meetings during the first reporting period. The second plenary meeting in May 2013 was held together with the first User and Advisory Board (UAB) workshop (part of the project management activities also included the organisation of the UAB workshop). This plenary meeting was focused on the requirements and use case elicitation, supported by comments from the UAB. The third plenary meeting in September 2013 was held in Helsinki at the headquarters of project partner Mirasys. This meeting focused on a visit to the Mirasys demo site, discussion of the use cases and focus areas in the demo, and on a discussion about the SECCRIT architectural framework and techno-legal aspects. In the subsequent, fourth plenary meeting in November 2013 in Valencia, we visited the local traffic control centre, which acts as the second demo-site for the SECCRIT project. In this meeting we also had a policy workshop and a discussion on the upcoming review meeting in February 2014.

WP1 activities in 2014 started with organising the first review meeting and a plenary meeting in February 2014 in Brussels. This lasted three days, and we discussed future plans and how to address reviewer comments. The next plenary meeting in June 2014 was held in Heidelberg at the premises of NEC. This plenary meeting focused on pending deliverables (especially D3.2 and D6.1). Also the feedback of the reviewers and the resulting To-Dos were discussed. Future dissemination activities were debated. The following plenary meeting was held together with the second User and Advisory Board (UAB) workshop (part of the project management activities also included the organisation of the UAB workshop). This plenary meeting took place in Vienna in September 2014. The UAB workshop was held as a joint event with EuroCloud Austria and the Federal Computing Centre of Austria – Bundesrechenzentrum (BRZ), the later kindly hosted the event. Cloud topics were presented in the following areas:

1. Commercial and industrial use of the Cloud
2. Cloud in the public administration and critical infrastructures
3. Security research activities regarding Cloud and critical infrastructures IT (SECCRIT)

The last plenary meeting in 2014 was held in Athens in November. In this meeting we discussed the feedback from OTE experts on our deliverables. It was also organised how RTD outputs are evaluated in the Demo (for D2.6). Some Dissemination and Exploitation activities (white paper, newsletter, med-1, UAB 2015, Conference workshop, ...) were discussed.

WP1 activities in 2015 started with organising the second review meeting and a plenary meeting in March 2015 in Valencia. This event lasted three days, and we discussed future plans. The next plenary meeting in June 2015 was held in Helsinki at the premises of Mirasys. This plenary meeting focused on the likely case studies and practical challenges involved in cloud adoption and operation. The consortium met once more in Vienna towards the end of the year to organise work on the demonstrators. As this was not an official plenary meeting, we don't include it in our list below.

The following document templates have been created for the consortium:

- Agenda
- Deliverable

- External Information
- Meeting Minutes
- Presentation
- Quarterly Management Report

The planning and maintenance of the project web site and IT infrastructure including the content management system Drupal (<https://www.seccrit.eu>), a twitter account (@SECCRIT), an e-mail list, and a Redmine-based (open-source project management web application) project portal including a SVN repository (<https://pp.seccrit.eu>) were also a part of the activities in WP1. The project portal is mainly used to record issues, organise meetings, provide wiki-pages for relevant parts, and as the project repository. The usage of all of these tools and facilities is intended to support our dissemination activities.

For engaging UAB members we have created Lol templates, SECCRIT introduction presentations, and an e-mail list. Public deliverables and papers as well as presentations from UAB workshops are made available to board members.

Monthly phone conferences for the Project Management Board to coordinate and manage the progress as well as quarterly steering committee meetings have been held – the latter always to organise the next plenary meeting.

The Project Management Board and partners typically meet every first Monday of the month via a phone conference:

- Status Update Tel.Co. 04/03/2013
- Status Update Tel.Co. 08/04/2013
- Status Update Tel.Co. 03/06/2013
- Status Update Tel.Co. 05/08/2013
- Status Update Tel.Co. 07/10/2013
- Status Update Tel.Co. 04/11/2013
- Status Update Tel.Co. 03/02/2014
- Status Update Tel.Co. 3/02/2014
- Status Update Tel.Co. 6/04/2014
- Status Update Tel.Co. 5/05/2014
- Status Update Tel.Co. 7/07/2014
- Status Update Tel.Co. 4/08/2014
- Status Update Tel.Co. 1/09/2014
- Status Update Tel.Co. 3/11/2014
- Status Update Tel.Co. 12/01/2015
- Status Update Tel.Co. 02/02/2015
- Status Update Tel.Co. 30/03/2015
- Status Update Tel.Co. 40/05/2015
- Status Update Tel.Co. 07/07/2015
- Status Update Tel.Co. 03/08/2015
- Status Update Tel.Co. 07/09/2015
- Status Update Tel.Co. 05/10/2015
- Status Update Tel.Co. 02/11/2015
- Status Update Tel.Co. 09/12/2015

Plenary Meetings (physical meetings):

- 01/2013 Vienna, AIT, Kick Off Meeting
- 05/2013 Vienna, BRZ & AIT, Plenary & UAB Meeting
- 09/2013 Helsinki, MIRASYS, Plenary Meeting & Demo-Site Visit
- 11/2013 Valencia, ETRA, Plenary Meeting & Demo-Site Visit
- 02/2014 Brussels, Plenary Meeting & Review Meeting
- 06/2014 Heidelberg, Plenary Meeting & Experts Talk
- 09/2014 Vienna, Plenary Meeting & UAB & EuroCloud Event
- 11/2014 Athens, Plenary Meeting
- 03/2015 Valencia, Review & Plenary Meeting
- 06/2015 Helsinki, Plenary Meeting & Demo-Site Visit

Steering Committee Meetings:

- 05/04/2013 Steering Committee Tel.Co.
- 23/07/2013 Steering Committee Tel.Co.
- 18/11/2013 Steering Committee Tel.Co.
- 25/02/2014 Steering Committee Meeting, Brussels
- 21/05/2014 Steering Committee Tel.Co.
- 23/01/2015 Steering Committee Tel.Co.

Initial resource problems of project partner OTE were resolved in M6, and options were discussed about how OTE could contribute valuable work in a more compact time than originally planned – this was mainly done in the valuable form of providing practice-based content for WP2 deliverables, and appropriate reviews.

As part of the project management activities, the leadership of WP3 was swapped with WP4 (ULANC <> IESE). This was agreed because of a change in Research Associate staff at ULANC which resulted in a change of expertise. The partners agreed on this bi-laterally and the consortium accepted the decision at the “Status Update Telco 03/06/2013”. The Project Officer was informed about this via e-mail and had no objection.

During the fourth plenary meeting, the consortium agreed that an amendment to the grant agreement should be created, which includes the Italian branch of Amaris Group as a linked third party. This amendment was sent off in June 2016. The REA had no objections about it and the amendment was granted.

Pertti Woitsch left Mirasys and was replaced by Petri Backstrom. Pertti Woitsch will continue to collaborate with the SECCRIT consortium as a valuable member of the User and Advisory Board

An NDA for cooperation with Med-1 (Israel) was created. The WP5 leader from NEC, Dr. Marcus Schöller, has left the company and his follower has been nominated. The reviewer comments have been answered and in response to it a scientific advisory board was established which consists of three people: Prof. Marcus Schöller (University of Reutlingen), Prof. Burkhard Schafer (University of Edinburgh) and Dr. Marcus Brunner (Swisscom). Their CVs can be found in the Annex, Section 5.1.

As well as continuing the project without interruption after the project's LEPPI, Dr. Frank Pallas, left KIT (replaced by Silvia Balaban) and after the coordinator, Dr. Markus G. Tauber, left AIT in 2015. In the latter case, continuation was established by Dr. Tauber who agreed on continuing

coordinating SECCRIT via part time employment at AIT. Contracts have been consolidated accordingly, and they have been confirmed by the PO.

4 Description of the potential impact, the main dissemination activities and the exploitation of results

WP 7 involved various dissemination activities and engagements with the UAB and other related projects, e.g. CUMULUS and Cloud4Europe, to address reviewer comments. AIT has organised the Second User and Advisory Board Workshop – Colleagues from both projects were present at his workshop.

Task T7.4 (Exploitation plan) started with the kick-off meeting in the last week of November 2013. The involved partners (especially those from industry) made an agreement on how to derive exploitation plans, and the aim has already been established, based on a set of tools (Market analysis, SWOT, etc.), to explore and derive a first version of the exploitation plans by the end of the second period.

In October 2014, Fraunhofer IESE received the Innovation Prize of the European Association of Research and Technology Organisation EARTO for their IND²UCE (Integrated Distributed Data Usage Control) framework. The prize was given in recognition of projects and research results that have the potential of triggering social or economic change.

Furthermore to address review comments, lectures were given to students as part of a lecture series¹² at the University of Applied Science in Eisenstadt (Austria) in September 2014.

Some of the SECCRIT partners participated in the organising committee of a seminar in the prestigious Dagstuhl Seminar Series¹³ on “Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations” in April 2015.

In total we have produced 31 deliverables. The seven output clusters have been evaluated in ten test cases. We have produced 36 peer-reviewed scientific papers of which 24 are collaborative. We have organised four User and Advisory Board workshops together with other EU initiatives and research projects, and we have organised a Dagstuhl seminar on “Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations”. Moreover, within the SECCRIT context, 11 student theses have been completed, and the SECCRIT project outputs have contributed to several lectures. Partners have succeeded to launch five follow-up projects to continue work on the individual output clusters, and various commercial exploitations are planned and vivid interaction with the standardisation community has been established.

4.1 Socio-economic impact and the wider societal implications

The wider impact of the project is threefold, via the vivid interaction with the User and Advisory Board the applicability of the research outputs to industrial needs has been assured and confirmed via our standardisation contributions. This in particular applies to the strong link between legal and technical aspects of the project outputs. Via the academic engagement in conferences and journal publications as well as in scientific events such as our Dagstuhl Seminar during which established scholars and scientists have been working on future interdisciplinary research agenda motivated by our work in SECCRIT. Last but not least, we have achieved educational impact by involving our outputs and findings in various lectures and training of students on BSc, MSc and PhD level.

¹² www.bitsthatbyte.at

¹³ <http://www.dagstuhl.de/15151>

4.2 User and Advisory Board & Liaisons with other Projects

As a result of the work carried out in task T7.3, liaison activities were established with Archistar, PRECYSE, CUMULUS, EuroCloud and the CLOSER conference where SECCRIT had a booth and where we participated in the European project space as academic partner. As part of task T5.1, a liaison and researcher exchange has been set up with colleagues from the CUMULUS project, from which a researcher will visit AIT to work together on assurance topics and to avoid redundancy in outcomes.

The third SECCRIT User and Advisory Board workshop was themed “The Future of Cloud”¹⁴. It took place on the 17th of June 2015. The event was hosted by AIT and organised together with EuroCloud Austria, presenting research from SECCRIT and its liaison projects PRISMACLOUD, CREDENTIALS, and CI2C.

A final engagement with members of the User and Advisory Board and with industrial stakeholders took place at the INFOCOM¹⁵ event in Athens on the 24th of November 2015. INFOCOM 2015 was the annual, major exhibition-congress with international impact where results, demonstrations, and other actions by Telecom providers, ICT vendors as well as many actors from the academic and research sector and from the European telecommunications market were presented. The SECCRIT project workshop was organized by OTE co-located with the INFOCOM event. The progress as well as the legal and RTD outputs, demos and exploitation plans were presented in relation to the legal framework created by the EU regarding Critical Infrastructures.

Earlier UAB workshops have been organised together with BRZ and other EURITAS members who represent European public administration / governmental data centre operators.

As a result of the work carried out, liaison activities have already been established with

- **Archistar**, a research project focusing on virtual cloud storage systems: It was agreed to pro-actively inform each other about research outputs.
- **PRECYSE**, an FP7 research project focused on cyber-security in CI: A total of three SECCRIT partners are participating in it (AIT, ETRA, and Valencia) the project has already identified some areas where synergies will be explored, especially the risk assessment methodologies and the demonstrations in the traffic management centre in Valencia.
- **CUMULUS**, an FP7 research project which focuses on assurance: A researcher from the CUMULUS project (Maria Krotsiani) spent some time at AIT to support the assurance activities in order to avoid redundancies in the anticipated outputs. A result of this collaboration was the following publication:
 - o Aleksandar Hudic, Maria Krotsiani, Markus Tauber, George Spanoudakis and Andreas Mauthe. A Multi-Layer and Multi-Tenant Cloud Assurance Evaluation Methodology. In: IEEE CloudCom 2014.
- **EuroCloud**, an organisation offering audits and certificates for cloud infrastructure providers: It was decided to inform each other about relevant outputs and to give SECCRIT the possibility to address the members of EuroCloud with surveys. [We

¹⁴ www.thefutureofcloud.org

¹⁵ <http://www.info-com.gr/en/#>

successfully organised a joint event with the EuroCloud Brunch in September 2014 together with our User and Advisory Board workshop.]

- **CLOSER**, the 4th International Conference on Cloud Computing and Services Science: CLOSER lists SECCRIT as academic partner.
- **CREDENTIAL**: The EU Project CREDENTIAL (Horizon 2020 program duration 10/2015-9/2018) engages in developing, testing, and showcasing innovative cloud-based services for storing, managing, and sharing digital identity information and other highly critical personal data with a demonstrably higher level of security and privacy than other current solutions. This is achieved by advancing novel cryptographic technologies and improving strong authentication mechanisms. The SECCRIT work on assurance is going to be extended and applied in the specific application domain of federated identity and access management in the cloud.
- **PRISMACLOUD**: The EU Project PRISMACLOUD (Horizon 2020 programme) addresses challenges in trustworthy cloud computing and yields a portfolio of novel agile cryptographic technologies to build security and privacy preserving services in the cloud. Its goal is to protect data throughout the whole lifecycle in the cloud and from end to end, i.e., it allows for trustworthy services on semi-trusted infrastructures. The work from SECCRIT plays an important role for PRISMACLOUD, and risk assessment results as well as the released guidelines and reference architecture will help to strengthen the project outcome.

4.3 Standardisation activities

The SECCRIT project's standardization activities are focusing on three SDOs: IETF, ONF, and ETSI. The details are given in the following paragraphs. In addition, ENISA is one of our user and advisory board members, and they are clearly an important dissemination target for SECCRIT outputs, for example in the areas of cloud security, risk assessment and resilience including metrics.

In the IETF, SECCRIT partners KIT, ULANC, and NEC have been monitoring working groups relevant to the work being conducted in SECCRIT. Since the IETF mainly standardises protocols, but not architectures, the work done in this standardisation body was seen as a possible source for standards that can be used in the SECCRIT solutions. One example is the standardised SYSLOG protocol that was used in the Tools for Audit Trails output cluster. In addition, it was observed that SECCRIT addresses relevant issues and partially fills gaps not yet addressed by standardisation activities, such as the ongoing work on Network Virtualization Overlays or Service Function Chaining.

Similarly, the work being conducted in the Open Networking Foundation (ONF) was monitored. Here, the work of the security group was of special interest. While the basic premises and issues identified in this group align well with the goals formulated for SECCRIT, the more detailed work lately, i.e., the testing of security properties of an SDN system and the hardening of the SDN control protocols and controller interfaces, is somewhat orthogonal to the work being conducted in SECCRIT.

Finally, SECCRIT partners have been active in the ETSI NFV ISG, in particular in the Reliability working group. This resulted in contributing to the still ongoing work on reliability estimation. More specifically, SECCRIT provided as input the architectural considerations and identified processes resulting from the Resilience Framework, in particular the Deployment Function. The contribution includes a discussion of deploying resilience patterns based on our reliability models of the

infrastructure, as well as service descriptions capturing the approach taken with the Deployment Function in SECCRIT. Our SECCRIT input has been accepted and will be included in the stable draft of this work item, expected to be published by early 2016.

More details about these activities can be found in D7.3 “Dissemination report year 3”.

4.4 Dissemination Activities

A web site has been launched, and social network profile at Twitter was created and was maintained on a regular basis.

This section provides a brief summary of dissemination activities carried out during the project. Complete details can be found in D7.3 “Dissemination report year 3”.

These are the figures achieved in the third periodic report:

- **38** peer reviewed **scientific papers** were accepted (at the point of compiling this report); **24 of them are joint** papers authored by multiple SECCRIT consortium partners.
- **25 presentations** of the SECCRIT project at different events
- **4 SECCRIT User and Advisory Board (UAB) workshops:** One meeting in Vienna (Austria) was hosted by AIT and organized together with EuroCloud Austria, presenting research from SECCRIT and its liaison projects PRISMACLOUD, CREDENTIALS, and CI2C. In addition, a final engagement with members of the User and Advisory Board and with industrial stakeholders took place at the **INFOCOM Event** in Athens (Greece). Two earlier ones with BRZ and EuroCloud in collaboration
- **11 student theses** successfully completed
- **6 newsletters** have been published to the User and Advisory Board
- A **White Paper** update was issued in November 2015
- **2 different lectures** about aspects of SECCRIT findings were given at the University of Applied Sciences Burgenland (Austria) and at Lancaster University.
- **An international Dagstuhl Seminar** entitled “Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations” was led by SECCRIT partners, and held in April 2015.

4.5 Exploitation activities

The SECCRIT Consortium identified the key exploitable results from the project, clearly marking those that have the chance to become potential products. Information can be found in D7.5 “Exploitation plan 2”. This deliverable also provides a brief market analysis per output cluster to support project members in finding appropriate exploitation channels and to coordinate the optimal exploitation of SECCRIT results.

Details about exploitable results, novelty highlights, identification of roles, bases for products, and further research work have been addressed. A comparison of SECCRIT RTD (Research and Technical Development) outputs and existing similar solutions in research and in the commercial

domain has been tackled. Furthermore, follow-up projects for the research outputs have been identified, and an initial agreement on IPR has been established.

In the context of the exploitation plan, the first step was to identify, from the project outcomes, those that can be exploited by the consortium once the project will be ended. This section focuses on the identification of the exploitable results. Below an overview table has been created to summarise all the exploitable results and the main characteristics of each one. It includes the timetable for their commercial use and registers if a result has associated patents, IPR ongoing issues or licenses or if there are any foreseen activities after the end of the project. Finally it identifies the owner of each result.

As the detailed description of each one of the project clusters has been provided in section 2, section 4.5.1 is focused on a very brief but useful description of the results to be exploited in the format of specific tables. Each table summarises the main functionality of the result, its purpose and the innovation that provides, the partners involved, their roles and activities. Furthermore, it is interesting to briefly describe how the result is foreseen to be exploited, as well as the additional deployment work that would be required. In the case of IPR protection, where measures are foreseen that are – different from those foreseen in SECCRIT Consortium Agreement -, they can be mentioned. Finally, the commercial contact for the exploitation of the result is also included in each table. As a result, sections 4.5.1 and 0 provide an overall picture of the exploitable results and their main characteristics.

4.5.1 Overview of exploitable products

Exploitable product	Type of exploitable foreground	Confidential	Sector(s) of application	Timetable, commercial or any other use	Patents or other IPR exploitation (licenses)	Owner and Other Beneficiary(s) involved
IND ² UCE Framework Policy Specification, Decision and Enforcement	Commercial exploitation of R&D results	Partially Closed Source	Data Privacy Protection, Cloud Management, Policy Specification	TRL 4 TRL 5 (Q1/2016)	Proprietary Software License (e.g, Software License for Permanent and Gratuitously Use ¹⁶)	IESE (owner)
Legal Guidance	General advancement of knowledge	No	Law	Currently on TRL3	Published in Journals	KIT
Deployment Function	Exploitation of R&D results via standards	No	Cloud Orchestration	Currently on TRL4, expected to increase in following years	N/A	NEC (owner)
Efficient and reliable storage for Audit Trails (Hybris), part of Tools for Audit Trails and Root Cause Analysis	General advancement of knowledge and commercial exploitation	No	Cloud Storage	Currently on TRL4, expected to increase in following years	IPR (underlying PoWerStore protocols)	NEC (owner)
Alert Monitoring Tool (Traffic and Transportation Security)	Commercial Exploitation of R&D results	No	Traffic Management and Transportation systems hosted in cloud	Currently TRL5, expected to arrive at TRL7 in two years' time	ETRA proprietary solution	ETRA (Owner), Ayto. Valencia (user)
Resilience framework including Anomaly Detection in the cloud	General advancement of knowledge	No	Cloud Management	Currently on TRL4, expected to increase in following years	N/A	ULANC



Independent Transparency-as-a-Service-Framework (<i>CloudInspector</i>), part of Tools for Audit Trails and Root Cause Analysis	General advancement of knowledge; maybe Commercial Exploitation of R&D results	No	Infrastructure-as-a-Service Providers and their tenants	Currently TRL5, expected to arrive at TRL9 within the next five years	Shared Source Software License, IPR	KIT
Consultancy supporting Risk Assessment Approach and Security Guideline	Commercial Exploitation of R&D results	No	Cloudification of critical infrastructure IT Services	TRL6, - related software is in support of consultancy services	N/A	AIT
Assurance Methodology & Framework	General advancement of knowledge	No	Cloud Management	TRL2	N/A	AIT
Video Surveillance System as a Service	Commercial Exploitation of R&D results	No	Security and surveillance service providers	TRL7 system prototype demonstration in pilot customer environment	MIRASYS proprietary solution	MIRASYS (owner)

TABLE 5: OVERVIEW OF EXPLOITABLE PRODUCTS

4.5.2 Details of exploitable project results

Exploitable result	IND ² UCE Framework
Functionality	Policy specification (PAP component), policy decision making (PDP component), policy enforcement (PEP and PXP components), policy management (PMP component), policy information resolution (PIP component)
Purpose	The IND ² UCE frameworks enables the flexible enforcement of data usage control policies. It empowers stakeholders to specify their security demands, which are technically enforced by the framework components.
Innovation	• Dynamic enforcement on multiple abstraction layers (cloud management and service layer) • Dynamic runtime behaviour through management component • Component-based approach allows easy integration of additional components (also during runtime) • Policies can include contextual information (e.g., resource usage on physical host) • PAP component allows the generation of different policy specification interfaces • Policy specification methodology allows the systematic elicitation of security policy templates within an application domain
Partner(s) involved	Fraunhofer IESE
Role and activities	Fraunhofer IESE has competencies in the area of data usage control and transfers the concept to other application areas and technologies. Fraunhofer IESE developed and maintains the IND ² UCE framework that is a technological solution to enable data usage control. In addition, several prototypical enforcement and execution points for different application areas such as VMware or HBase as well as a policy information point for VMware are available.
How the result will be exploited	Fraunhofer IESE is in contact with different companies interested in the area of data usage control. The exploitation is done in three steps: • Presentation of data usage control concept including research prototypes such as VMware enforcement. Presentation of IND²UCE lab at Fraunhofer IESE in Kaiserslautern. • Data usage control workshop to identify the potential of data usage control enforcement based on specific application scenarios of the companies. • Proof of concept to evaluate the technological feasibility and integration of the IND²UCE framework into the existing software system
Additional research and development work	Fraunhofer IESE is doing additional research in the area of data usage control for mobile devices and how to enable data usage control in smart ecosystems. In addition, Fraunhofer IESE investigates the trade-off between security (special focus on data usage control) and user experience.
IPR protection measures	Software License, Contract Research



Commercial contacts	Fraunhofer IESE, IND2UCE@iese.fraunhofer.de Christian Jung, christian.jung@iese.fraunhofer.de , +49 631 6800 2146
Target groups	Companies that need a flexible framework for policy enforcement and companies that are focusing on data-driven business models.

Exploitable result	Legal Guidance
Functionality	Legal guidance of data protection and evidence law concerning cloud computing
Purpose	Permit to identify gaps within the legal framework or ways of addressing these issues
Innovation	Technologies developed within SECCRIT permit to fulfil the legally requested requirements
Partner(s) involved	KIT
Role and activities	Identified the relevant legal issues with regard to cloud computing
How the result will be exploited	Academic education, PhD theses, scientific articles, follow-up research projects
Additional research and development work	The SECCRIT findings will be further extended on the KASTEL project.
IPR protection measures	Knowledge and results protected through publications in journals
Commercial contacts	KIT, Center for Applied Legal Studies (ZAR)
Target groups	Students, relevant associations, regulatory bodies, scientific audience

Exploitable result	Deployment Function
Functionality	Instantiating resilience patterns based on cloud user and infrastructure provider input (cf. D4.2)
Purpose	Automated support of critical components that use resilience patterns to increase reliability
Innovation	Includes reliability aspect in structure-aware placement of components, developed generic method to communicate placement constraints to OpenStack
Partner(s) involved	NEC
Role and activities	- Developed concept of reliability-aware placement function taking a service description and infrastructure info into account - Implemented proof-of-concept for active-backup resilience pattern for OpenStack

How the result will be exploited	Gained knowledge has been used to provide input to ETSI NFV standardization
Additional research and development work	Research and development in the areas of automated generation of infrastructure input, support for a wider range of resilience patterns and possibly different cloud management systems are needed
IPR protection measures	No IPR, concept and all algorithms have been published
Commercial contacts	NEC Laboratories Europe Simon Oechsner (simon.oechsner@neclab.eu)
Target groups	NEC business units/subsidiaries: Telecom Carrier Business Unit, NetCracker

Exploitable result	Efficient and reliable storage for Audit Trails (Hybris)
Functionality	Interconnects generator of audit trail information with the storage provided by Hybris
Purpose	Provides efficient and reliable cloud storage for audit trail information as provided by the CloudInspector
Innovation	Improves the efficiency of write-back-based byzantine fault tolerant storage. The efficiency gain is achieved by combining lightweight cryptography, erasure coding and metadata write-backs, where readers write-back only metadata to achieve strong consistency.
Partner(s) involved	NEC
Role and activities	- Developed and evaluated storage protocol based on metadata write-back - Implemented Hybris SYSLOG interface towards audit trail generator
How the result will be exploited	Integration into cloud solutions and products, such as for example the HYDRASor series.
Additional research and development work	Further research is required related to storage integrity, key management and integration into existing solutions. This work will be conducted in further EU projects, e.g., H2020 TREDISEC
IPR protection measures	IPR on storage protocol, software released under LGPL v2.1 license
Commercial contacts	NEC Laboratories Europe Wenting Li (wenteng.li@neclab.eu)
Target groups	NEC business units: Business Innovation Unit, System Platform Business Unit



Exploitable result	Alert Monitoring Tool (Traffic and Transportation Security)
Functionality	The main functionality of this cloud-based service tool is to monitor several facilities (instances of traffic management and transportation system) with a unique user interface. It is able then to correlate security events at tenant level with events occurring at service layer
Purpose	To be a central point for tenant managers to monitor security properties, so operators of the infrastructure can have deep insights on what is happening in the infrastructure in a very quick way
Innovation	Correlation of security logs and events, with a clear differentiation at cloud infrastructure, tenant and service level
Partner(s) involved	ETRA is the owner of this solution, currently it has been customised for the Valencia Municipality, and makes use of the IND ² UCE Framework (Fraunhofer IESE) and AD Framework (ULANC)
Role and activities	The platform has been built according to ETRA needs and to satisfy one of their customers, Valencia Municipality, but the platform is expected to cover in the future other areas where ETRA has market presence like e.g. Energy
How the result will be exploited	The result will be mainly exploited by ETRA, it will be used to support internally the technical team responsible for the maintenance of the facilities. The platform, in addition, strengthens the ETRA security portfolio, since other competing traffic and transportation entities are not able to deliver services in the cloud with a similar security approach.
Additional research and development work	There is still work to do to meet the standards of a production environment, we believe at least two years of work will be needed to industrialise the solution. In any case part of the development will be further extend and work in the PRISMACLOUD project.
IPR protection measures	All IPR are proprietary of ETRA, except for the IND ² UCE framework which is under software license, contract research.
Commercial contacts	ETRA has a strong market presence especially with regards to traffic and transportation systems, in the case of traffic management systems ETRA counts on 55% of the market share. The developments produced in SECCRIT may impact all the current commercial contacts of the company
Target groups	ETRA technical support service, public administrations which are running ETRA systems, and future contracts in the areas of traffic and transportation

Exploitable result	Anomaly Detection
Functionality	The main function of the anomaly detection component is to detect anomalous pattern in the cloud infrastructure (cf. D4.2 and D4.3) This is important to support overall resilience to challenges.
Purpose	To provide real-time detection of the anomalies in cloud infrastructure to the subsequent stages of the overall resilience loop, i.e. D ² R ² +DR (cf. D4.2).



Innovation	The anomaly detection approach is designed to be memory-less and adaptive to new threats and challenges online and in real time under minimal computational cost. This overcomes the commonly used signature-based solutions that are currently dominating the domain of cloud security.
Partner(s) involved	ULANC
Role and activities	Develop, design and implement the anomaly detection technique for the overall resilience framework.
How the result will be exploited	Mainly through the scientific papers, follow-up research projects and teaching. The knowledge is used to develop anomaly evaluation framework which is part of the Advance Networking module (SCC 365) of the Master level course at Lancaster University.
Additional research and development work	Further research is required in online identification of anomalies to complement the detection phase and integration of detection and classification into cloud specific security solutions.
IPR protection measures	No IPR, concept has been published
Commercial contacts	Lancaster University (n.shirazi@lancaster.ac.uk)
Target groups	Students, research community and standardization and regulatory bodies.

Exploitable result	Independent Transparency-as-a-Service-Framework (<i>CloudInspector</i>)
Functionality	On-demand real-time auditing of current situation as well as evidence gathering for root-cause analysis. Both in case of Infrastructure-as-a-Service Clouds.
Purpose	Fulfilment of legal requirements data protection and civil law (D5.1, D2.7): Providing an independent view on the current state of the tenant's virtual resources within the cloud. Possibility to collect evidences for a later Root Cause Analysis in court.
Innovation	Overcome the Lack of Transparency by providing current situation of virtual machines on-demand of tenants and collecting information for Root Cause Analysis in court. In both cases used data is based on Cloud Infrastructure Management System independent audit sources of physical cloud nodes. In this way failures or misconfigurations of the Cloud Infrastructure Management System are detectable for tenants and may provable in court.
Partner(s) involved	KIT
Role and activities	• Development and design of <i>CloudInspector</i>, part of Tools for Audit Trails and Root Cause Analysis. • Implemented proof-of-concept of <i>CloudInspector</i> for OpenStack and VMware ESXi. • Implemented CloudInspector SYSLOG interface towards storage implementation • Evaluation of <i>CloudInspector</i> in SECCRIT use and testcases.



How the result will be exploited	Mainly through the scientific papers, follow-up research projects and teaching. Based on the results of the bachelor thesis regarding to figures on the market, it being under consideration to found a spin off. Source code will be available under a shared source license.
Additional research and development work	Further research and development is required. There is still work necessary to enable the use in a production environment, we believe at least five years of work will be needed to industrialise the solution. Additionally, enhancements to network and storage abstraction and further mechanisms of independence from Cloud Infrastructure Management System could be analysed, as well as supplementary legal requirements from data protection law.
IPR protection measures	Shared Source Software License ``
Commercial contacts	Roland Bless and Matthias Flittner Karlsruhe Institute of Technology Institute of Telematics Chair Prof. Zitterbart Zirkel 2 D-76131 Karlsruhe Tel.: +49 721 608-46400 Fax: +49 721 608-46789 E-Mail: telematics@tm.kit.edu
Target groups	Infrastructure-as-a-Service Providers and their tenants; Research Community

Exploitable result	Consultancy supporting Risk Assessment Approach and Security Guideline
Functionality	The risk assessment approach developed in SECCRIT (D3.1) can help customers managing risk when moving services to the cloud. Likewise the Security Guideline (D3.4) supports cloudification by extending software development life cycles for secure cloud service migration.
Purpose	Both outputs in combination support the cloudification process of IT services of critical infrastructure providers.

Innovation	There is currently no cloudification support including risk assessment and best practices guidelines with a special focus on critical infrastructure IT, see state of the art analysis in D3.1 & D3.4 and e.g. [HaTm+2015], [WHTP2015] and [PsTm+205] [HaTm+2015] Hudic A., Tauber M., Loruenser T., Krotsiani M., Spanoudakis G., Mauthe A., Weippl E., "A Multi-Layer and Multi-Tenant Cloud Assurance Evaluation Methodology", IEEE CloudCom 2014 [WHTP2015] Wagner C., Hudic A., Tauber M., Pallas F., "Impact of Critical Infrastructure Requirements on Service Migration Guidelines to the Cloud", IEEE Future Internet of Things and Cloud - IEEE FiCloud 2015 (to appear) [PsTm+205] Paudel S., Tauber M., Wagner C. and Ng W., "Categorization of Standards, Guidelines and Tools for Secure System Design for Critical Infrastructure IT in the Cloud", IEEE Cloud-CPS 2014 at IEEE CloudCom 2014,
Partner(s) involved	AIT; KIT (legal) and NEC
Role and activities	AIT lead the work on both D3.1 and D3.4 - D3.1 was supported by NEC and D3.4 by KIT-Legal (as well as by most technical partners for showing how their RTD output can support a secure cloud migration)
How the result will be exploited	Both, D3.1 and D3.4 will be used for consulting dedicated customers
Additional research and development work	Additional work is required to provide tool support and further consolidate vulnerability catalogues
IPR protection measures	None, results have been published.
Commercial contacts	Dr. Martin Stierle AIT Austrian Institute of Technology T: +43 50 550 4166 F: +43 50 550 4150 Donau-City-Strasse 1 1220 Wien martin.stierle (at) ait (ac.at)
Target groups	Individual customer requests

Exploitable result	Assurance Assessment Methodology and Framework
Functionality	Modelling and measurement of monitoring artefacts to confirm high level definition for security properties
Purpose	Providing an indication of an overall level of security without revealing underlying infrastructure details.

Innovation	Continuous assurance is provided by aggregation over information regarding individual constituent components of a cloud service. Representation via a bit mask allows a scalable aggregation and flexible policies by e.g. ordering in the bit mask or with logical operations are being used to aggregated the bitmasks per component. – see [HaTm+2015] [HaTm+2015] Hudic A., Tauber M., Loruenser T., Krotsiani M., Spanoudakis G., Mauthe A., Weippl E., "A Multi-Layer and Multi-Tenant Cloud Assurance Evaluation Methodology", IEEE CloudCom 2014
Partner(s) involved	AIT (supported by ULANC via academic papers)
Role and activities	AIT lead this activity
How the result will be exploited	The methodology will be further developed in future projects (e.g. CREDENTIAL, www.credential.eu)
Additional research and development work	The current set of security properties is limited to the special needs of the SECCRIT demonstrator; this needs to be extended. Scalable approaches are required to deploy monitoring points in infrastructure, tenant and service level. Dependencies between components need to be resolved automatically
IPR protection measures	None, results have been published.
Commercial contacts	Dr. Martin Stierle AIT Austrian Institute of Technology T: +43 50 550 4166 F: +43 50 550 4150 Donau-City-Strasse 1 1220 Wien martin.stierle (at) ait (ac.at)

Exploitable result	Video Surveillance System as a Service (VSaaS)
Functionality	Mirasys VMS services can be distributed within the networked infrastructure. Video recording and content analysing servers can be run at the edge (traditional onsite approach), servers/services can be centralized (public/private cloud approach) or servers can be partly at the edge and partly in the cloud. Cloud-based surveillance system architecture can be designed to be deployed on public cloud, private cloud, or a combination of both, which can be called hybrid cloud.
Purpose	Purpose is to provide security operators and/or end users an opportunity to purchase video management service with a fixed monthly fee, instead of buying hardware equipment, software licenses and installation services.
Innovation	VSaaS business model and enabling technologies are emerging on the markets.
Partner(s) involved	Mirasys is the owner of this solution

Role and activities	Testing of the VSaaS model and Mirasys VMS technology match has been made in proof-of-concept type of environments already with several pilot customers. Pilot customers are not part of the SECCRIT consortium.
How the result will be exploited	The results will be exploited commercially. The most promising part of the results, applicable for short term exploitation are the private clouds and networks.
Additional research and development work	Mirasys needs to carefully follow the market perception for cloud-based security services and operators. There is a risk that the market does not adopt the VSaaS model or adoption comes with a remarkable delay. The market seems to hesitate due to concerns about data leaking, privacy, etc. It needs to be studied if there are additional technical countermeasures to be taken in order to enable market adoption. Other areas are work-in-progress, and relate to such things as Mirasys VMS service-to-service communications protocols (being converted to be fully HTTPS-based for our next major software version release), or file storage system for video and audio storage for very large-scale and/or multi-tenant systems (also under development and being addressed in forthcoming software releases), as well as the need to migrate away from monolithic client applications to more light-weight user-interface applications where network services provide APIs to actual business logic
IPR protection measures	All IPR are proprietary of Mirasys.
Commercial contacts	Mirasys Ltd. is one of the leading suppliers of open platform Video Management Systems (VMS). More than 50.000 customers use Mirasys systems with nearly one million surveillance cameras connected. VSaaS prototype demonstration has been done in operational environment with a couple of pilot customers: 1. A mid-sized regional telecom operator in Western Finland. 2. A security operator start-up in Sweden. 3. A small data center operator in Poland. 4. A large security service provider in Southern Finland. Contact: Jouni Räihä, Principal Product Manager, Integrations, tel. +358 50 594 1517 Email: jouni.raiha(at)mirasys.com

Target groups

Mirasys is targeting for the Mirasys VMS software to become a suitable solution and service platform for service providers that are or wish to become providers of cloud-based video surveillance service providers. End users of these video surveillance service providers would benefit from using VSaaS because of its flexible technical features or because of its business model.

TABLE 6: EXPLOITABLE PRODUCTS

5 Annex

5.1 CVs of Scientific Advisory Board Members

5.1.1 Prof. Dr.-Ing. Marcus Schöller

Dr. Marcus Schöller joined the computer science department of Reutlingen University in September 2014 as a full professor for cloud computing. Before that, he was a senior researcher at NEC Laboratories Europe, Germany. He received his diploma in computer science in 2001 and his doctorate in engineering in 2006 on the topic of robustness and stability of programmable networks from University of Karlsruhe, Germany. Afterwards he held a postdoc position at Lancaster University, UK, focusing his research on autonomic networks and network resilience. Marcus has been working on multiple EU projects, e.g., ANA, ResumeNet, SECCRIT, with a focus on network function resilience and critical services on the cloud. Moreover, he was vice-chair of the ETSI's 'Reliability and Availability' WG within the Industry Specification Group on Network Function Virtualisation. His interests include security and dependability of networks, systems, and cloud environments, as well as privacy aspects of ICT, and future network architectures. He published his research results in multiple journal and conference papers, books, and patent applications.

5.1.2 Dr. Markus Brunner

Dr. Marcus Brunner is head of standardization in the strategy and innovation department of Swisscom. He received his Ph.D. from the Swiss Federal Institute of Technology (ETH Zurich) in 1999. He is active in research and standardization since 20 years with experience in programmability of networks and services, cloud technology for IT and network service providers, and automation of network and IT. For example, he represents Swisscom in various international organizations on software-defined networking, and the future telecommunication technologies like Network Function Virtualization (Nfv). One of the strategic focuses of Swisscom is getting secure cloud offerings for IT and telco workloads. For example, Swisscom joined the trusted computing group to standardize secure boot and security attestation of virtual machines.

5.1.3 Professor Burkhard Schafer

Professor Burkhard Schafer is Director and Co-founder of the SCRIPT centre for IT and IP law, and member of the management group of the RCUK funded CREATE research network that explores the future of copyright in the digital economy. He holds degrees in computer linguistics, logic and law from the Universities of Munich, Mainz and Lancaster. His main field of research is the interface between computer technology, science and the law. He has published more than 90 papers in the field of legal expert system design, the semantic web, and on legal responses to new technologies from a comparative perspective. He is Senior Visiting Research Fellow of the Centre for Social Innovation research at the University of New Brunswick, Canada; member of the Rechtsinformatik (Legal informatics) working group of the German Gesellschaft fuer Informatik and member of the executive of the International Association for AI and Law. He is Associate Editor (Law and AI) of Script-ed and Member of the Editorial Board of the Oxford Journal of Law, Probability and Risk; the Journal for Law and Information Technology and Jusletter IT. He has been PI on several externally funded grants, including Co-PI on the EU funded FF POIROT project on computer technologies for fraud detection, and member of the AEEC consortium, an EU AGIS funded project to analyse that state of regulation of digital evidence in Europe and to develop a standardised training programme in Forensic Computing and Digital Evidence for Judges, that was offered for the first time in 2009.

5.2 Logos

5.2.1 Project Logo



FIGURE 4: SECCRIT LOGO MONOCHROME



FIGURE 5: SECCRIT LOGO, COLOUR

5.2.2 Partner Logos



FIGURE 6: AIT AUSTRIAN INSTITUTE OF TECHNOLOGY GMBH



FIGURE 7: AMARIS TECHNOLOGIES GMBH & AMARIS TECHNOLOGIES SRL



FIGURE 8: ETRA INVESTIGACIÓN Y DESARROLLO, S.A.



FIGURE 9: FRAUNHOFER IESE



Karlsruhe Institute of Technology

FIGURE 10: KARLSRUHE INSTITUTE OF TECHNOLOGY



FIGURE 11: MIRASYS LTD.



FIGURE 12: NEC EUROPE LTD., UK




FIGURE 13: HELLENIC TELECOMMUNICATIONS ORGANIZATION S.A. (OTE)



FIGURE 14: LANCASTER UNIVERSITY



AJUNTAMENT DE VALENCIA

FIGURE 15: AJUNTAMENT DE VALENCIA