



## Secure Collaborative Development of Cloud Application Deployment Models

Vladimir Yussupov<sup>1</sup>, Michael Falkenthal<sup>1</sup>, Oliver Kopp<sup>2</sup>,  
Frank Leymann<sup>1</sup>, Michael Zimmermann<sup>1</sup>

<sup>1</sup>Institute of Architecture of Application Systems,

<sup>2</sup>Institute of Parallel and Distributed Systems,  
University of Stuttgart, Germany,

[lastname]@informatik.uni-stuttgart.de

---

BIBTEX

```
@inproceedings{Yussupov2018_SecureCollaborativeDeploymentModels,  
author    = {Yussupov, Vladimir and Falkenthal, Michael and Kopp, Oliver and  
            Leymann, Frank and Zimmermann, Michael},  
title     = {{Secure Collaborative Development of Cloud Application Deployment  
            Models}},  
booktitle = {Proceedings of The Twelfth International Conference on Emerging  
            Security Information, Systems and Technologies (SECURWARE  
            2018)},  
publisher = {Xpert Publishing Services},  
pages     = {48--57},  
month     = {September},  
year      = {2018}  
}
```

The full version of this publication has been presented at  
SECURWARE 2018.

<https://www.iaia.org/conferences2018/SECURWARE18.html>

© 2018 IARIA. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IARIA.



# Secure Collaborative Development of Cloud Application Deployment Models

Vladimir Yussupov\*, Michael Falkenthal\*, Oliver Kopp†, Frank Leymann\*, and Michael Zimmermann\*

\*Institute of Architecture of Application Systems

†Institute of Parallel and Distributed Systems

University of Stuttgart, Stuttgart, Germany

email: [lastname]@informatik.uni-stuttgart.de

**Abstract**—Industrial processes can benefit considerably from utilizing cloud applications that combine cross-domain knowledge from multiple involved partners. Often, development of such applications is not centralized, e.g., due to outsourcing, and lacks trust among involved participants. In addition, manual deployment of resulting applications is inefficient and error-prone. While deployment can be automated using existing modeling approaches, the issues of data confidentiality and integrity in exchanged deployment models have to be addressed. In this paper, we tackle security challenges posed by collaborative cloud application development. We present a policy-based approach for modeling of security requirements in deployment models. Furthermore, we propose a method of peer-to-peer model exchange that allows enforcing modeled requirements. To validate our approach we apply it to Topology and Orchestration Specification for Cloud Applications (TOSCA), an existing cloud applications modeling standard, and describe the prototypical implementation of our concepts in OpenTOSCA, an open source toolchain supporting TOSCA. Usage of the resulting prototype in the context of a described model exchange process allows modeling and enforcement of security requirements in collaborative development of deployment models. We then conclude the paper with a discussion on limitations of the approach and future research directions.

**Keywords**—Collaboration; Security Policy; Confidentiality; Integrity; Deployment Automation; TOSCA.

## I. INTRODUCTION

Modern computing paradigms have great potential for accelerating the 4<sup>th</sup> industrial revolution, often referred to as Industry 4.0 [1]. One notable example is the rapidly evolving field of cloud computing [2], which allows on-demand access to potentially unbounded number of computing resources. Combined together with ubiquitous sensors usage in the context of the Internet of Things (IoT) [3], cloud computing facilitates the development of composite, cross-domain applications tailored specifically for automation and optimization of manufacturing. The overall complexity of the development process, however, might become a significant obstacle for industries willing to benefit from cloud applications.

A typical cloud application today has a composite structure consisting of numerous interconnected and heterogeneous components [4]. Deploying such complexly-structured applications in a manual fashion is error-prone and inefficient. Therefore, various deployment automation approaches exist. One well-established automation technique relies on the concept of *deployment models* that specify application structure along with the necessary deployment information. Automated processing of such models considerably reduces the deployment’s complexity and minimizes required efforts. Another significant benefit, which improves the portability and reusability aspects of the application development process, is that standardized models can be exchanged instead of separate application components.

One common cloud application development scenario in the context of Industry 4.0 is a collaboration [5] among several multidisciplinary partners responsible for separate parts of the application [6]. The final goal of this collaboration is to combine all parts into a complete and deployable cloud application. Collaborative development can significantly benefit from the portability and reusability properties of deployment models. However, since not all parties are known in advance, e.g., due to task outsourcing or changes in organizational structure, the issues of intellectual property protection in decentralized settings arise. For instance, confidential information like sensor measurements and proprietary algorithms might be subject to various *security requirements*, including protection from unauthorized access and verification of its integrity. Therefore, modeling and enforcement of such requirements aimed at specific parts of deployment models, have to be supported.

In this work, we focus on the aspects of secure collaborative development of cloud applications’ deployment models. The contribution of this paper is a method for modeling and enforcement of security requirements in deployment models which combines the ideas of sticky policies [7], policy-based cryptography [8], and Cryptographic Access Control (CAC) [9]. We describe how security requirements aimed at data protection in modeled cloud applications can be expressed using dedicated security policy types and analyze which parts of deployment models need to support the attachment of security policies. As a next step, we elaborate on how modeled security requirements can be enforced in a peer-to-peer exchange of deployment models. To validate our concepts, we apply them to an existing OASIS standard called Topology and Orchestration Specification for Cloud Applications (TOSCA) [10], [11], which specifies an extensible, provider-agnostic cloud modeling language [12]. As a proof of concept, we describe the prototypical implementation of the presented concepts in OpenTOSCA [13], an opensource ecosystem for modeling and execution of TOSCA-compliant deployment models. The resulting prototype used in the context of the proposed decentralized model exchange serves as a means to model the discussed security requirements and enforce them along the model’s exchange path. Finally, we discuss the limitations of our approach and describe possible improvements.

The remainder of this paper is structured as follows. We describe the fundamentals underlying this work in Section II and discuss a motivational scenario in Section III. In Section IV, we present concepts for modeling and enforcement of security requirements in collaborative deployment models development. In Section V, we apply the concepts to a TOSCA-based deployment modeling process. The details about the prototypical implementation in OpenTOSCA are discussed in Section VI. In Section VII, we describe related work and Section VIII summarizes this paper and outlines future research directions.

## II. FUNDAMENTALS

In this section, we provide an overview of several important concepts which serve as a basis for our work, namely: (i) deployment automation of cloud applications by means of deployment modeling approaches, (ii) usage of policies as means to specify non-functional system requirements, (iii) and a brief coverage of access control mechanisms.

### A. Deployment Modeling

The compound application structure and increased integration complexity make it non-trivial to automate the deployment of modern cloud applications [4]. The concept of deployment modeling aims to tackle the automation problem, and there are several known approaches including imperative and declarative modeling [4], [14], [15]. Both paradigms are based on the idea of creating a description, or deployment model, sufficient enough for deploying a chosen application in an automated fashion. What makes these modeling approaches different is the way how corresponding deployment models are implemented.

In case of the declarative modeling [14], a deployment model is a structural model that conveys the desired state and structure of the application. Essential parts of the declarative deployment model include a specification of application's components with respective dependencies and necessary connectivity details. As a result, the model might contain binaries or scripts responsible for running some application's components, e.g., a specific version of Apache Tomcat, or a predefined Shell script for running a set of configuration commands. In addition, a description of non-functional system requirements in some form can be included into the model. Some examples supporting this type of modeling include Chef [16] and Juju [17] automation tools, as well as TOSCA. This type of models relies on the concept of deployment engines, which are able to interpret a provided description and infer a sequence of steps required for successful deployment of the modeled application.

Compared to declarative approach, the imperative modeling [14] focuses on a procedure which leads to automatic application deployment. More specifically, an imperative model describes (i) a set of activities corresponding to the required deployment tasks which need to be executed, (ii) the control and data flow between those activities. One robust technique for this modeling style is to use a process engine, e.g., supporting standards like Business Process Execution Language (BPEL) [18] or Business Process Model and Notation (BPMN) [19], that can execute provided imperative models in an automated fashion.

A combination of declarative and imperative approaches is also possible. In general, creating both types of models requires efforts from the modeler. However, the imperative modeling approach is generally more time-consuming and error-prone, since multiple heterogeneous components need to be properly orchestrated. Moreover, the structure of the application might change frequently which requires to modify imperative models. To minimize required modeling efforts, imperative models might be derived from the provided declarative models [4].

One important aspect of deployment models is that apart from valid descriptions they also need to include various files related to described software components and other parts of the application, e.g., scripts, binaries, documentation and license details. As a result, the term deployment model usually refers to a combination of all the corresponding metadata and application files required for automatically deploying a target application.

### B. Policies

One well-known approach [20] for separation of non-functional requirements from the actual functionalities of a target system relies on the usage of policies. Essentially, a *policy* is a semi-structured representation of a certain management goal [21]. The term management here is rather broad, as it might refer to different aspects of management, e.g., high-level corporate goals or more low-level, technology-oriented management goals. For instance, from the system's perspective, performance, configuration, and security are among the classes of non-functional requirements that can be described using policies. Additionally, various policy specification languages exist in order to simplify the process of describing such requirements in a standardized manner [20]. From the high-level view, policies only declare the requirements which then have to be enforced using dedicated enforcement mechanisms [22].

The idea to specify security requirements in policies dates back to at least the 1970s [20]. Depending on the level of details security policies might specify, e.g., privacy requirements for the whole system or for particular data objects. In information exchange scenarios, security policies specified on the level of data objects have to be ensured during the whole exchange process [23]. For this reason, all receivers have to be aware of specified policies and enforcement must happen, e.g., by means of globally-available security mechanisms. Similarly, deployment models in collaborative application development are constantly exchanged and parts of them might be subjects to security policies. So-called *sticky policies* [23] is an approach to propagate policies with the data they target. This approach can be combined with cryptography in order to ensure that data is accessed only when requirements specified in policies are satisfied. Multiple approaches to combine sticky policies with different cryptographic techniques such as public key encryption or Attribute-Based Encryption (ABE) exist [24].

### C. Access Control

A secure information system must prevent disclosure (confidentiality) or modification (integrity) of sensitive data to an unauthorized party and ensure that data are accessible (availability) [22]. These requirements can be enforced by assuring only authorized access to the system and its resources. Commonly, this process is referred to as *access control* and there exist multiple well-established access control mechanisms. For example, in Discretionary Access Control (DAC) mechanism, the access is defined based on the user's identity. This results in access rules that are specified specifically for this identity, e.g., in the form of an access control matrix [25]. Another well-known access control mechanism is called Role-Based Access Control (RBAC) where access is granted or denied based on the user roles and access rules defined for these roles.

One disadvantage of aforementioned access control mechanisms is that they commonly rely on some centralized trusted authority, making it difficult to implement them in large scale and open systems [9]. The idea of CAC is based on well-known cryptographic mechanisms and regulates access permissions based on the possession of encryption keys. In CAC, the stored data are encrypted and can only be accessed by those users who have the corresponding keys. One benefit of this approach is that the data owner can grant keys to receivers of his choice using established key distribution mechanisms, thus enforcing the access control without relying on the trusted third party.

### III. MOTIVATIONAL SCENARIO

Developing distributed cloud applications and analytics applications in the context of Industry 4.0 typically requires combining numerous heterogeneous software components [26], [27]. Commonly, this process implies a collaboration among experts from various domains, such as data scientists, infrastructure integrators, and application providers. Furthermore, resulting applications are often required to be deployable on demand and, thus, are expected to be in the form of deployment models that allow automating application provisioning [6], [28].

An example of a collaborative cloud application development depicted in Figure 1 involves four participants responsible for distinct parts of the application. When joined together, all developed parts of the application, e.g., software components, datasets, and connectivity information, comprise a complete and provisioning-ready deployment model. In this scenario, the main beneficiary who orders the application from a set of partners and has exclusive rights on the resulting deployment model is called the *Application Owner*. The *Infrastructure Modeler* is responsible for integrating different components, such as analytics runtime environments, databases, or application servers. Moreover, two additional co-modelers are involved in the development process, namely a *Data Scientist* and a *Dataset Provider*. The former develops a certain proprietary algorithm, whereas the latter provides a private dataset, e.g., comprised of sensor measurements obtained from a combination of various cyber-physical systems used in production processes.

In contrast to the *Application Owner* who has full rights on the resulting deployment model, other participants might be subjects to security restrictions with respect to certain application parts. For example, access to the dataset provided by the *Dataset Provider* might need to be restricted to some of the involved parties. Similarly, the *Data Scientist* might want to impose a certain set of security requirements on the provided algorithm. Since the final infrastructure must include all corresponding sub-parts that were provided directly or indirectly by other participants, the *Infrastructure Modeler* is responsible for preparation and shipping of the finalized deployment model to the *Application Owner* who is then able to create new instances of the application on demand.

Generally, collaborative processes from various fields share some common characteristics. For instance, according to Wang et al. [29] such issues as (i) a *dynamically changing sets of participants*, (ii) the *lack of centralization*, (iii) *intellectual property and trust management issues*, and (iv) *heterogeneity of exchanged data* are important in collaborative development of computer-aided design models. Likewise, the lack of knowledge about all participants involved in collaborative cloud application development makes it difficult to establish a centralized interaction among them. Possible reasons include outsourcing of development tasks and introduction of additional participants due to rearrangements in organizational structures. Since no strict centralization is possible, communication with known participants happens in a peer-to-peer manner. Another important aspect of collaborative cloud application development is its iterative nature. Since exchanged deployment models might be impartial or require several rounds of refinement, a potentially complicated sequence of exchange steps is possible for obtaining a final result. Therefore, deployment models need to be exchanged in collaborations in a way that simplifies the overall process and enforces potential security requirements.

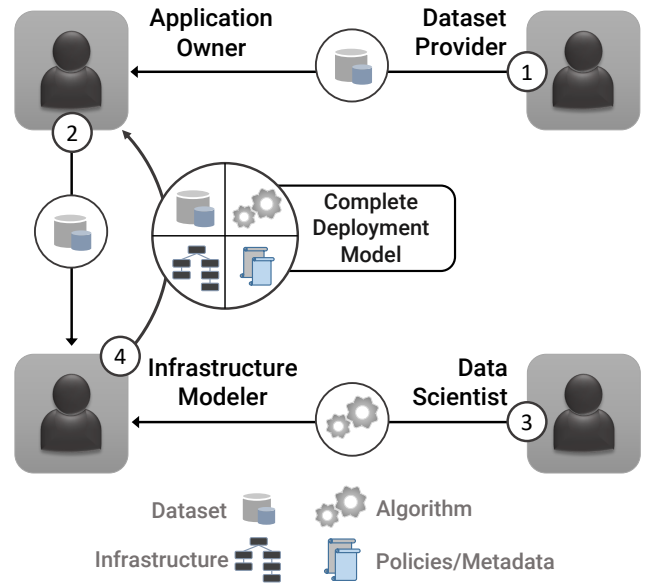


Figure 1. A collaborative application development scenario.

A deployment model, generally, can be exchanged either in a self-contained form or on a per-participant basis. In the former case, the deployment model is self-contained and its content is the same for all participants, whereas in the latter case its content is fragmented according to some rules separately for each participant. Sometimes, however, exchanging deployment models on a per-participant basis interferes with the actual goals of the collaboration. For example, in exchange sequence shown in Figure 1 the dataset is firstly passed directly to the *Application Owner* by the *Dataset Provider*. For integration of the dataset into the final model, the *Infrastructure Modeler* needs to model the required infrastructure, e.g., a Database Management System (DBMS) and related tooling. As only the *Application Owner* has full rights on all parts of the application, the provided dataset has to be protected from unauthorized access. Intellectual property issues become even more complex in highly-dynamic scenarios when multiple parties continuously exchange partially-completed deployment models. Unfortunately, encrypting an entire deployment model does not solve the problem since models might be intended to remain partially-accessible by parties with limited access rights. Apart from confidentiality problems, the authenticity and integrity of passed deployment models and their parts might be subjects to verification requirements. For instance, the *Application Owner* might need to check if an algorithm was actually provided by the *Data Scientist* and no changes were made by other parties. In such case, signing the hash value of an entire deployment model is not suitable as integrity of individual model's parts have to be verified. Hence, it should be possible to verify distinct parts of deployment models independently.

The aforementioned scenario highlights several important issues in collaborative development of deployment models which need to be solved, namely (i) confidentiality, authenticity, and integrity requirements of each involved participant have to be reflected in the model, (ii) various levels of granularity for these requirements need to be considered: from full models to its separate parts, and (iii) a method to enforce modeled requirements in a peer-to-peer model exchange is needed.

#### IV. MODELING AND ENFORCEMENT OF SECURITY REQUIREMENTS

Intellectual property in collaborations has to be protected from both, external and internal adversaries with respect to their relation to the process. The former describes any attacker from outside of the collaboration, i.e., who is not participating and is not reflected in any kind of agreements, e.g., Service Level Agreements (SLAs). Conversely, the latter refers to a dishonest party involved in the process. We focus on internal adversaries and data protection issues involving known parties.

This section presents an approach to ensure the fulfillment of security requirements in collaborative development of deployment models. Our approach relies on the well-established concept of representing non-functional requirements via policies [30], [31], [32], [33]. The semantics of security requirements is analyzed to derive a set of action and grouping policies. The former type represents cryptographic operations allowing to enforce confidentiality and integrity requirements, inspired by the idea of policy-based cryptography [8]. The latter type simplifies grouping parts of models which are subjects to action policies. Both policy types are data-centric and attachment happens with respect to a certain entity or a group of entities in the manner of sticky policies [23] to preserve the self-containment property of deployment models. The access control enforcement is inspired by the idea of CAC [9].

##### A. Assumptions

To focus on internal adversaries, we assume that participants establish bidirectional secure communication channels for data exchange and that the modeling environment of every involved participant is secure. We employ an “honest but curious” [34], [35], [36] adversary model in which adversaries are interested in reading the data, but avoid modifications to remain undetected. Despite the absence of modifications made by adversaries, authenticity and integrity requirements still need to be modeled and enforced. For instance, participants might want to track changes or verify the origin of some specific part in the model.

When describing how data encryption can be modeled, we assume that no double encryption is needed for distinct parts of deployment models. We do not distinguish between read and write rights when discussing access control based on cryptographic key possession. Therefore, a participant with the required key is assumed to have full access rights on the corresponding entity. For efficiency reasons, we adopt symmetric encryption for ensuring the confidentiality of data.

##### B. Security Policies in Collaborative Deployment Models

An assumption that data is exchanged in a secure manner among the participants does not guarantee that all involved parties can be trusted. Therefore, security requirements are important even under the secure communication channels assumption. Security requirements we focus on are: (i) protection of data confidentiality in deployment models, and (ii) verification of data integrity and authenticity of deployment models. On the conceptual level, two distinct types of policies, namely *encryption policy* and *signing policy*, can be distinguished. The former is aimed to solve the confidentiality problem, whereas the latter targets integrity-related requirements. However, having a completely encrypted deployment model does not solve the confidentiality problem, since a party with limited rights will not be able to access the parts of the application which were

intended to remain accessible. Similar problem might arise for a signature of the complete packaged deployment model, e.g., in a form of an archive, since it will not be possible to check what exactly was changed unless all files are also signed separately as a part of the process. More specifically, if only the hash of an entire deployment model was signed, there will be no way to distinguish which specific part of the model is invalid. Therefore, we need to model security policies on the level of atomic entities in deployment models to support collaborations similar to the scenario described in Section III.

Naturally, if only parts of deployment models are subjects to confidentiality requirements, enforcement of encryption and signing policies must affect only respective entities. In our approach, an encryption policy attached to a certain entity of the deployment model signals that it has to be encrypted. In a similar manner, if a certain entity of the deployment model needs to be signed, the corresponding signing policy needs to be linked with it. In both cases, policies represent actual keys that are going to be used for encryption or signing. Since not all collaborations can rely on a centralized way to manage policies, the deployment model has to be transferred together with corresponding policies attached to its entities. The keys bound to policies, however, cannot be embedded, as deployment models will no longer remain suitable for sharing with all possible participants in a self-contained fashion. In such cases, either participants with proper access control rights can receive such models, or the models have to be split on a per-participant basis. Since not all scenarios favor participant-wise model splitting, a policy needs to be linked with a specific key in a decoupled manner to preserve self-containment of a deployment model. As a side effect of decoupling keys from policies, existing key distribution channels can be utilized independently from deployment model exchange channels.

For linking policies with particular keys, we need to maintain unique identifiers for every key involved in the collaboration. Since not all participants know each other, one simple solution is to compute a digest of the key and use it as an identifier or additionally combine it with several other parameters such as algorithm details, participant identifier, etc. Another option is to use identifiers which include some partner-specific parts so that policies can be easily identified. Several important points have to be mentioned here. Linking the policy only with the unique key identifier is not enough for decryption since the modeler needs to know the algorithm details to perform decryption. Such information can be provided either as properties of a policy itself or be a part of the key exchange. Additionally, specifically for encryption there is no obvious way to distinguish if the policy was already applied and the data is in encrypted state when a deployment model is received. Although the data format after encryption will not be identical to the original entity’s format, checking this difference for every modeled entity is not efficient. For this reason, a policy needs to have an attribute stating that it was applied. Due to the usage of symmetric encryption, generating a respective *decryption policy* is unnecessary as it is identical to the encryption policy.

Conversely, the verification of signing policies differs from the encryption process since private keys are used for signing and certificate chains of one or more certificates containing the public key and identity information are used for verification. As a result, there are two options: to follow the encryption approach and decouple certificates from policies, or, to embed

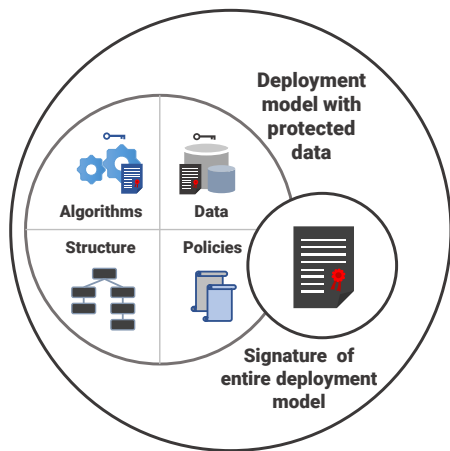


Figure 2. A conceptual model of the signed deployment model.

certificates into policies to simplify the verification process. While certificates are meant for distribution, there is one caveat in the embedding of certificates approach, however. Certificates commonly have a validity period and verification must be able to deal with the cases when certificates embedded into policies are no longer valid. Since such verification is more an issue of a proper tooling, the certificates are embedded into policies.

Unlike file artifacts, e.g., software components or datasets, which are referenced from models and supplied alongside with them, some sensitive information, e.g., model's properties, might be directly embedded into models. For instance, if user credentials for a third-party service have to be passed from one modeler to another and no other participant is allowed to see them, then these properties must be encrypted. Sometimes such properties also need to be verified, e.g., the Service Owner might want to check if the endpoint information for a third-party service was actually modeled by the Infrastructure Modeler. Therefore, an additional caveat one has to consider is that not only distinct artifacts, but also separate parts of artifacts might require encryption or signing. The corresponding artifact in this case has to store these properties with the modeled security requirements being enforced, e.g., encrypted or signed.

Hence, we need two more policy types: *encryption grouping policy* and *signing grouping policy* which contain lists of properties within an artifact that have to be encrypted or signed, respectively. From the conceptual point of view, the discussed policies can be classified as *action* and *grouping* policies. The former includes policies representing an action, i.e., encryption or signing, whereas the latter identifies groups of entities which require the action. As a result, the corresponding grouping policies are linked with the desired action policies, i.e. with actual keys which will be applied to selected properties.

### C. Integrity and Self-Containment of Deployment Models

When security policies are modeled and enforced, the resulting deployment model contains a combination of encrypted and signed artifacts and properties. Integrity check at this point allows to verify the state of modifications and authenticity of entities modeled by other participants. However, verification of the entire deployment model's integrity including modeled security policies and other attached metadata requires an additional signature on the level of deployment model.

For this purpose we adopt the technique analogous to signing of Java archives (JARs) [37]. Essentially, a packaged deployment model is some sort of an archive containing grouped artifacts. It is then possible to assume the presence of a meta file similar to manifest in JARs, which provides the list of all contents plus some additional information. In situations when such manifest file does not exist, it can easily be generated by traversing the contents of a corresponding deployment model.

As both, integrity of the model's parts that are targeted by security requirements and integrity of the entire deployment model have to be considered, an enhanced packaging format is needed. The enhanced structure of a deployment model consists of its original content as well as the content's signature files. The latter is achieved via a combination of: (i) a manifest file with digests for every file, (ii) a signature file consisting of digests for every digest given in the manifest file plus the digest of the manifest file itself, and (iii) a signature block file consisting of a signature generated by the modeler and the certificate details. The resulting conceptual model is shown in Figure 2. To make a signed deployment model distinguishable from regular deployment models, the signature has to be generated in a standardized fashion, e.g., it can be stored in a predefined folder inside the package or entire deployment models can be archived along with the generated signature information.

One important issue is that, technically, there is no fixed concept of a deployment model in collaboration. Since parts of cloud applications might be exchanged separately or merged together, the definition of the exchanged deployment model is changing throughout the process. Thus, it is mandatory to preserve the self-containment of modeled security requirements on the level of atomic entities. Firstly, security policies are always included to the deployment model since they are tightly-coupled with target entities. With respect to actual entities, the problem is trivial in case of encryption since locations of files or properties remain unchanged and only their state changes. In other words, whether the encrypted entity is exported from or imported into the modeling environment, the information about encryption is always available. Conversely, signatures of modeled entities have to be created as separate files since embedding them might not always work. For instance, embedding a signature into the application's source code might result in an incorrect behavior at runtime. This leads to a requirement of generating and storing signatures in a self-contained manner when signing policies enforcement happens.

In contrast, the signature of an entire deployment model reflects a snapshot of its state at a particular point in time, e.g., when the deployment model was packaged by a certain participant. Semantically, this signature does not mean that all content of the deployment model belongs to a signing party, but only captures the state of the deployment model at export time. In our approach, we use this external signature only for integrity verification at import time, but do not explicitly store it if verification was successful. However, if stored in a centralized or decentralized manner, this type of signature might form an expressive log of all export states which can later be utilized for audit and compliance checking purposes.

### D. Enforcement of Security Policies

As participants of collaboration might not know all involved parties, every side has to maintain a set of permissions for known participants, e.g., in a form similar to the access matrix




Figure 3. Actions of a collaborating participant.

model [22]. In our case, permissions have to reflect which policies are available to which participant and are therefore used for export and distribution of keys. One caveat is that in long sequences of steps there will be cases when a party does not know which rights with respect to the specific key have to be defined for some of the involved parties. The rules in such collaborations rely on various types of agreements, such as SLAs, which define the lists of trusted parties. Hence, we handle only explicitly mentioned access rights defined by participants and forbid transitive trust [38] propagation.

To enforce security policies in collaborations, participants have to follow a set of actions shown in Figure 3. A new or existing deployment model can be imported into the participant’s modeling environment. Signatures are verified for an existing deployment model before import. An entire model’s signature is verified first and if verification is successful, all signed entities are verified next. If certificate chains are embedded, all certificates must be valid. The import is aborted in case some signatures or certificates are invalid. Participant might request keys needed for encrypted entities and if access is granted by the key owner, keys can be imported into the modeling environment and used for decryption. The policy enforcement at export time happens transparently for participants as entities always get encrypted if the respective keys are present. Since decryption is only possible when the key is available, the encryption at export is ensured by the modeling environment.

Afterwards, participants can model additional security requirements and export a modified deployment model. One issue related to signatures and mutual modifications of the same entity is whether to keep the obsolete signature information. Since the original content of the entity has to be modified, we consider it being a new entity which can be modeled separately eliminating the problem of handling several signatures altogether. At export time, all modeled requirements are enforced with respect to the keys available in modeling environment. The decrypted data get encrypted again, in case the corresponding key is present and the entity was decrypted previously. Only signatures modeled by the participant who performs the export are generated. All entities that were signed by others remain in a self-contained state after import and thus exported in a regular fashion.

Generated signatures must be linked with corresponding modeling constructs. For instance, for every signed file the corresponding signature files must be added as additional linked references, e.g., following a predefined name format “filename#sigtype.sig”. Signing properties requires a slightly different approach. Since properties are parts of artifacts and are subject to certain policies, their signatures have to be grouped with respect to the policy. This results in generation of the

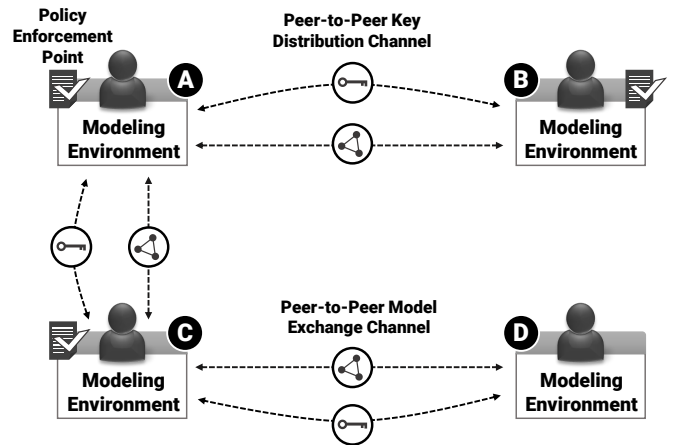


Figure 4. Model and key exchange in collaborations.

combined signature file and linking it with the artifact which holds the signed properties. Signature of this file is, again, generated similar to JAR files signing, but in this case the generated artifact contains the details about signed properties.

Figure 4 shows communication infrastructure for collaboration described in Section III. As key distribution is decoupled from the model exchange, two peer-to-peer channel types are distinguished. Generally, not all participants need to communicate with each other. For example, in outsourcing case, a contractor grants rights to the ordering party based on the contract rules and does not need to communicate with others. Therefore, access permissions of the ordering party have to also reflect access rules for the part of deployment model provided by the contractor. The access to encrypted data is inquired by requesting a key using the corresponding policy identifier. Without having a centralized Policy Enforcement Point (PEP) [39], [40], every participant’s modeling environment acts as a separate PEP which regulates access control permissions based on inter-participant agreements. Participants are responsible for maintaining proper access control permissions including transitive cases.

## V. STANDARDS-BASED SECURE COLLABORATIVE DEVELOPMENT OF DEPLOYMENT MODELS

In this section we discuss the specifics of collaborative development of deployment models using TOSCA. We analyze which TOSCA modeling constructs might require protection and describe how our concepts can be applied to this technology.

### A. TOSCA Application Model

TOSCA [10] is a cloud application modeling standard which allows to automate the deployment and management of applications. The structure of a TOSCA application is characterized by descriptions of application’s components with corresponding connectivity information, modeled as a directed, attributed graph which is not necessarily connected. In TOSCA terminology the entire application model is called a *Service Template*, whereas the connectivity information is a subpart of it and referred to as a *Topology Template*. The management information in TOSCA terms is called *Management Plans*. This information is necessary for execution and management of applications throughout their lifecycle and can be represented, e.g., in a form of BPEL [18] or BPMN [19] models.

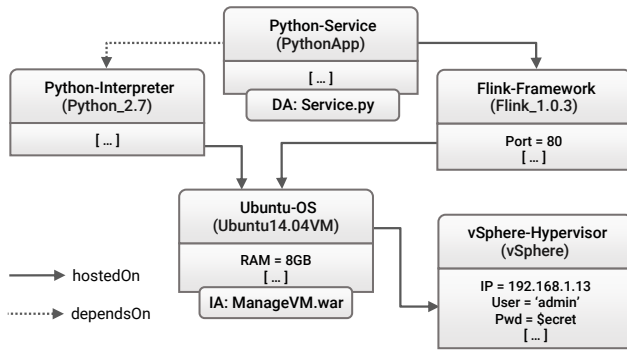


Figure 5. A simplified TOSCA model of a cloud application.

A simplified TOSCA topology of a Python cloud service [6] is shown in Figure 5. It consists of several *nodes* representing software components which are connected with directed edges describing the *relationships* among them. TOSCA differentiates between *entity types* and *entity templates*, where the term entity might refer to distinct TOSCA entities such as nodes, relationships, artifacts, or policies. Such separation eases reusing modeled TOSCA entities, since the semantics is always defined in the corresponding type. For instance, the “vSphere” in Figure 5 is called a *Node Type*. It describes a generic setup of a vSphere virtualization platform and defines all required configuration properties. Apart from defining common properties, any Node Type might provide definitions of interface operations required for managing its instances. For example, a virtual machine node might have two interface management operations, namely “start” and “stop” implemented using Java web services. Correspondingly, the “vSphere-Hypervisor” represents a particular instance of the “vSphere” *Node Type* and in TOSCA terms is referred to as a *Node Template*.

For deployment and management of the cloud service, all required artifacts have to be modeled, e.g., the application files and implementations of management interface operations. The artifact entity in TOSCA can be of two types, namely deployment artifacts (DA) and implementation artifacts (IA). The former defines an executable required for materialization of a node instance. The latter is a representation of an executable which implements a certain interface management operation.

One of the main goals of deployment models is to make cloud applications portable and reusable. For this reason TOSCA introduces a self-contained packaging format called Cloud Service Archive (CSAR). Essentially, it is an archive containing all application-related data necessary for automated deployment and management, including, e.g., the model definitions, artifact files, policies and other metadata. In addition, it contains a TOSCA.meta file which describes files in the archive similarly to a manifest file in JARs.

### B. Security Requirements for TOSCA Entities

Several TOSCA modeling constructs can be associated with confidential information or be subjects to integrity checks. Modeled application files, i.e., artifacts in TOSCA terms, is one obvious example. All artifacts are always modeled as Artifact Templates of particular Artifact Type in TOSCA, e.g., a Java web application artifact is a template of Web application Archive (WAR) Artifact Type. While Artifact Type is a generic

entity which does not store any sensitive data, the Artifact Templates include actual application files. However, in TOSCA specification there is no standard way to describe security requirements using policies for Artifact Templates. To provide such modeling capabilities, an extension to TOSCA is needed. Since properties are defined at the level of Types in TOSCA, e.g., Node Types, it is useful to have a mechanism allowing to enforce security requirements at this level. Semantically, this would mean that encryption or signing policies have to be applied to all Node Templates of a certain Node Type. TOSCA does not offer a standard way of attaching policies to specific properties, thus a proper way to enforce protection of properties at the level of Node Types is needed as well.

### C. TOSCA Policy Extensions

To support the attachment of security policies to aforementioned TOSCA entities we introduce several extension points. All policies are defined in a dedicated extension element which belongs to a chosen entity. A simplified XML snippet in Figure 6 shows extension policies for Artifact Templates and Node Types from Figure 5. For Artifact Templates, a security policy is attached in a separate element directly to the Artifact Template. Essentially, an Artifact Template is a container grouping related files in a form of file references. We treat Artifact Templates as atomic entities meaning that policies are applied to all referenced files which makes the semantics of modeled security requirements clearer. If some referenced files need to be distributed without enforcement of policies, they can be modeled as separate Artifact Templates.

A combination of two policy types has to be defined in a dedicated extension element for encryption and signing of properties. A modeler has to specify a list of property names that must be encrypted or signed as well as to attach a corresponding action policy. These extensions allow participants to model desired security requirements for parts of the CSAR.

```
<ArtifactTemplate name="Python-Service" ...>
  <Policies>
    <Policy applied="false" name="encryption"
      policyType="csar:EncryptionPolicyType"
      policyRef="csar1:c0e9a0e7".../>
  </Policies>
  <ArtifactReferences>
    <ArtifactReference ref=".../Service.py"/>
  </ArtifactReferences>
</ArtifactTemplate>
...
<NodeType name="vSphere" ...>
  <PropertiesDefinition>...</PropertiesDefinition>
  <Policies>
    <Policy ... name="signing" .../>
    <Policy ... name="signedprops" .../>
  </Policies>
</NodeType>
```

Figure 6. Example of TOSCA extension policies specification in XML.

The introduced extensions, however, do not offer modeling capabilities for signing the entire CSAR. These two notions of integrity might contradict with each other, since a party having parts of the cloud service belonging to other parties is required to sign them as well. Hence, we separate the integrity check for a specific part of the model from an integrity check of the entire CSAR leaving the latter outside of TOSCA modeling.



The Policy Types and Templates representing action and grouping policies are lightweight. The Encryption Policy Type defines a key's hash value, an algorithm, and a key's size as its properties. In the corresponding Policy Template, these properties are populated using the respective key's data. Similarly, the Signing Policy Type has public key's hash and related certificate chain as its properties, filled in using the given key. Certificate chain can be embedded, e.g., in a form of a Privacy Enhanced Mail (PEM) encoded string in case of X509 [41] certificates. The only property defined in grouping policies is a space-separated list of property names. This Policy Type is abstract and is not directly bound to any specific entity. Therefore, the tooling is responsible for checking the consistency of specified property names in attached policies.

#### D. Self-Contained CSAR

Preservation of CSAR's self-containment property after enforcement of modeled policies requires embedding the signature information for artifacts and properties into the corresponding entities. More specifically, when a signature for an artifact is created, it has to be placed along with other files referenced in the artifact. For the signature of properties, one artifact containing all properties' signatures needs to be generated and attached to the corresponding Node Template. Following this approach, modeled entities remain self-contained even in case they are being reused in other Service Templates.

## VI. PROTOTYPICAL IMPLEMENTATION

In this section we describe the prototypical implementation of the presented concepts. The prototype is based on the OpenTOSCA ecosystem, an open source toolchain for development and execution of TOSCA-compliant cloud applications. The OpenTOSCA ecosystem consists of such tools as Winery [42], [43], OpenTOSCA Container [13], and Vinothek [44].

Winery is the core part for implementation of the presented concepts, as most of them are coupled with the modeling process. Winery is a feature-rich modeling environment for TOSCA-compliant applications. It is written in Java programming language and uses Angular for the frontend. The prototype is open source and available via Github [45]. As discussed in Section IV, in our approach every modeler is required to use a local Winery instance due to the absence of a centralized environment. Since keys are used for enforcement of policies, Winery is extended to support key management functionalities. This includes storing, deletion, and generation of symmetric and asymmetric keys. For key storage we rely on usage of Java's Java Cryptography Extension KeyStore (JCEKS) keystore for storing all imported keys together. Assuming that Winery runs in a local and secure environment of a distinct party, publishing keys is not problematic since keys never leave the modeler's environment. This approach, however, has to be extended to support multiple-owner Winery instances. Corresponding policies are generated based on selected keys. For key distribution, a partner-wise specification of access control lists for security policies is added to Winery. Every participant needs to maintain the list of partner-specific rules negotiated by means of agreements in collaborations. Therefore, whenever a key is requested by some party, the key access rights are defined based on the local rules in Winery. All functionalities are accessible via the corresponding REST endpoints.

The prototype supports modeling of security requirements via Winery's built-in XML editors for respective TOSCA entities. Winery stores modeled TOSCA entities in a decoupled manner making a concept of CSAR important only at export or import time. At import time, CSARs are disassembled into distinct entities to prevent storing duplicates. In a similar manner, at export time CSARs are assembled from all the entities included in the chosen Service Template. This results in an issue that TOSCA meta files are not explicitly stored and are generated on-the-fly. Enforcement of modeled security policies at export time for selected TOSCA entities, e.g., Service Templates or Artifact Templates, happens in case specified keys are present in the system. Signatures for files in Artifact Templates are generated as additional files in the same Artifact Template. If the files of Artifact Templates are subjects to both, encryption and signing requirements, then the signatures of plain and encrypted files are attached. This allows verifying the integrity of target files to both, authorized and unauthorized parties. Signatures for properties are grouped as a separate Artifact Template of type "Signature" which is attached to the respective Node Template. This ensures the self-containment property of deployment models. If policies were applied, the corresponding attribute is set to signify this fact. After encryption and signing requirements are enforced, an external signature of a CSAR is generated using a so-called master key, which is specified by the modeler for the whole environment as discussed in Section IV. The corresponding certificate or chain of certificates for this external signature is embedded into the CSAR and is used for verification at import time. This signature is verified first at import time and is not stored if verification succeeds, since the CSAR is decomposed into distinct separately-stored entities. Import does not happen in case if integrity checks were not successful. In case keys requested by a modeler were provided, they can be imported and used for decryption of entities. Finally, only the modeler who has an entire set of keys is able to decrypt and deploy the final application. Deployment and execution in OpenTOSCA Container then happens in a regular manner, since the CSAR contains the original deployment model.

## VII. RELATED WORK

The problem of data protection in outsourcing and collaboration scenarios appears in works related to different fields. Multiple works attempt to tackle security-related problems using centralized approaches. Wang et al. [29] present a method for protecting the models in collaborative computer-aided design (CAD), which extends RBAC mechanism by adding notions of scheduling and value-adding activity to roles. Authors propose to selectively share data to prevent reverse engineering. However, no clear description how to enforce the proposed model is given. Cera et al. [46] introduce another RBAC-based data protection approach in collaborative design of 3D CAD models. Models are split into separate parts based on specified role-based security requirements to provide personalized views using a centralized access control mechanism. Li et al. [30] propose a security policy meta-model and the framework for securing big data on the level of Infrastructure as a Service (IaaS) cloud delivery model using sticky policies concept. Policies are loosely-coupled with the data and the framework relies on a trusted party which combines policy and key management functionalities and enforces the access control. Huang et al. [47] introduce a

set of measures allowing to protect patients data in portable electronic health records (EHRs). Authors propose a centralized system which combines de-identification, encryption, and digital signatures as means to achieve data privacy. Li et al. [34] describe an approach based on the Attribute-Based Encryption which helps to protect patient's personal health records in the cloud. In this approach, data is encrypted using keys that are generated based on the owner-selected set of attributes and then published to the cloud. Users can only access the data in case they possess corresponding attributes, e.g., profession or organization. More specifically, users are divided into several security domains and the attributes for these domains are managed by corresponding attribute authorities. Decryption keys, therefore, can be generated independently from data owners by the respective attribute authorities.

A number of approaches focus on the data encryption in outsourcing scenarios. Miklau and Suciu [48] introduce an encryption framework for protecting XML data published on the Internet. Contributions of the work include a policy specification language available in the form of queries and a model allowing to encrypt single XML documents. Access control is enforced based on key possession. Vimercati and Foresti [49] discuss fragmentation-based approaches for protecting outsourced relational data. The authors elaborate on several techniques allowing to split up the given data based on some constraints into one or more fragments and store them in a way to protect confidentiality and privacy. For instance, data can be split into two parts and stored on non-communicating servers. Whenever constraints cannot be satisfied for some attributes, the encryption is used. In the follow-up work, Vimercati et al. [50] present a way to enforce selective access control using the cryptography-based policies. Authors propose to use key derivation mechanisms to simplify the distribution of keys.

To the best of our knowledge, none of the discussed approaches successfully tackles our problem of deployment models protection in collaborative application development scenarios. Most of the discussed approaches rely on the idea of a trusted party which can regulate the access control. While it is desirable to have a central authority, in many cases it is unrealistic, leading to a need for peer-to-peer solutions. Moreover, having focus only on separate security requirements like encryption or strong assumptions about the underlying data make these approaches not suitable for the described problems.

## VIII. CONCLUSION AND FUTURE WORK

In this work, we showed how security requirements can be modeled and enforced in collaborative development of deployment models. We identified sensitive parts in deployment models and proposed a method which allows protecting them based on a combination of existing research work. For validation of the presented concepts, we applied them to TOSCA, an existing OASIS standard, which specifies a provider-agnostic cloud modeling language. The resulting prototypical implementation is based on the modeling environment called Winery, which is a part of the OpenTOSCA ecosystem, an open source collection of applications supporting TOSCA.

One issue in our approach that has to be optimized is the way keys are distributed. We rely on the fact, that not all participants need to exchange keys which, however, does not solve the scalability problem. If  $N$  keys were used for encryption, eventually all of them will be used in key distribution. For

improving the efficiency, the key derivation techniques, e.g., described by Vimercati et al. [50], can be used to reduce the number of keys that need to be exchanged. Another problem for future work is the generalization of the adversary model. Since deployment models can be intentionally corrupted by an adversary, there is a strong need to store the provenance information which describes deployment model's states at every export with respect to certain collaboration. Having such provenance information stored in some accessible form makes it possible to track the entire collaboration history with all the deployment model states that were existing throughout the process. For this reason, one might employ a centralized system, which will also simplify the policy enforcement and key distribution processes, or store the provenance in a decentralized fashion, e.g., by utilizing the blockchain technology [51].

Finally, there is a pitfall for cases when files are modeled in a form of references, e.g., if they reside on a remote server. Encrypting and signing such files completely changes the verification semantics as only the references are checked. This is not safe since the actual content behind the reference can be changed multiple times by the data owner without changing the reference itself. Moreover, the usage of references invalidates the self-containment property of deployment models. In the future work, referenced files need to be materialized at export time which solves this problem and preserves deployment models in a self-contained state.

## ACKNOWLEDGMENT

This work is funded by the BMWi projects *SePiA.Pro* (01MD16013F) and *SmartOrchestra* (01MD16001F).

## REFERENCES

- [1] M. Hermann, T. Pentek, and B. Otto, "Design principles for industrie 4.0 scenarios," in 2016 49<sup>th</sup> Hawaii International Conference on System Sciences (HICSS). IEEE, 2016, pp. 3928–3937.
- [2] P. M. Mell and T. Grance, "Sp 800-145. the NIST definition of cloud computing," Gaithersburg, MD, United States, Tech. Rep., 2011.
- [3] L. Atzori, A. Iera, and G. Morabito, "The internet of things: A survey," *Computer networks*, vol. 54, no. 15, 2010, pp. 2787–2805.
- [4] U. Breitenbücher et al., "Combining declarative and imperative cloud application provisioning based on toasca," in Proceedings of the IEEE International Conference on Cloud Engineering (IEEE IC2E 2014). IEEE Computer Society, March 2014, pp. 87–96.
- [5] T. Kvan, "Collaborative design: what is it?" *Automation in construction*, vol. 9, no. 4, 2000, pp. 409–415.
- [6] M. Zimmermann, U. Breitenbücher, M. Falkenthal, F. Leymann, and K. Saatkamp, "Standards-based function shipping – how to use toasca for shipping and executing data analytics software in remote manufacturing environments," in Proceedings of the 2017 IEEE 21<sup>st</sup> International Enterprise Distributed Object Computing Conference (EDOC 2017). IEEE Computer Society, 2017, pp. 50–60.
- [7] G. Karjoth, M. Schunter, and M. Waidner, "Platform for enterprise privacy practices: Privacy-enabled management of customer data," in International Workshop on Privacy Enhancing Technologies. Springer, 2002, pp. 69–84.
- [8] W. Bagga and R. Molva, "Policy-based cryptography and applications," in International Conference on Financial Cryptography and Data Security. Springer, 2005, pp. 72–87.
- [9] A. Harrington and C. Jensen, "Cryptographic access control in a distributed file system," in Proceedings of the 8<sup>th</sup> ACM symposium on Access control models and technologies. ACM, 2003, pp. 158–165.
- [10] OASIS, Topology and Orchestration Specification for Cloud Applications (TOSCA) Version 1.0, Organization for the Advancement of Structured Information Standards (OASIS), 2013.

- [11] OASIS, Topology and Orchestration Specification for Cloud Applications (TOSCA) Primer Version 1.0, Organization for the Advancement of Structured Information Standards (OASIS), 2013.
- [12] A. Bergmayr et al., “A systematic review of cloud modeling languages,” *ACM Comput. Surv.*, vol. 51, no. 1, Feb. 2018, pp. 22:1–22:38.
- [13] T. Binz et al., “Opentosca – a runtime for toasca-based cloud applications,” in *Service-Oriented Computing*. Berlin, Heidelberg: Springer, 2013, pp. 692–695.
- [14] C. Endres et al., “Declarative vs. imperative: Two modeling patterns for the automated deployment of applications,” in *Proceedings of the 9<sup>th</sup> International Conference on Pervasive Patterns and Applications*. Xpert Publishing Services (XPS), Feb. 2017, pp. 22–27.
- [15] U. Breitenbücher, K. Képes, F. Leymann, and M. Wurster, “Declarative vs. imperative: How to model the automated deployment of iot applications?” in *Proceedings of the 11<sup>th</sup> Advanced Summer School on Service Oriented Computing*. IBM Research Division, Nov. 2017, pp. 18–27.
- [16] Chef. [Online]. Available: <https://www.chef.io/> [retrieved: July, 2018]
- [17] Juju. [Online]. Available: <https://jujucharms.com/> [retrieved: July, 2018]
- [18] OASIS, Web Services Business Process Execution Language (WS-BPEL) Version 2.0, Organization for the Advancement of Structured Information Standards (OASIS), 2007.
- [19] OMG, Business Process Model and Notation (BPMN) Version 2.0, Object Management Group (OMG), 2011.
- [20] R. Boutaba and I. Aib, “Policy-based management: A historical perspective,” *Journal of Network and Systems Management*, vol. 15, no. 4, Dec 2007, pp. 447–480.
- [21] R. Wies, “Using a classification of management policies for policy specification and policy transformation,” in *Integrated Network Management IV*. Springer, 1995, pp. 44–56.
- [22] P. Samarati and S. C. di Vimercati, “Access control: Policies, models, and mechanisms,” in *International School on Foundations of Security Analysis and Design*. Springer, 2000, pp. 137–196.
- [23] S. Pearson and M. Casassa-Mont, “Sticky policies: An approach for managing privacy across multiple parties,” *Computer*, vol. 44, no. 9, 2011, pp. 60–68.
- [24] Q. Tang, *On Using Encryption Techniques to Enhance Sticky Policies Enforcement*, ser. CTIT Technical Report Series. Netherlands: Centre for Telematics and Information Technology (CTIT), 2008, no. W6TUG-31/TR-CTIT-08-64.
- [25] B. W. Lampson, “Protection,” *ACM SIGOPS Operating Systems Review*, vol. 8, no. 1, 1974, pp. 18–24.
- [26] M. Falkenthal et al., “Opentosca for the 4<sup>th</sup> industrial revolution: Automating the provisioning of analytics tools based on apache flink,” in *Proceedings of the 6<sup>th</sup> International Conference on the Internet of Things*, ser. IoT’16. New York, NY, USA: ACM, 2016, pp. 179–180.
- [27] T. Binz, U. Breitenbücher, O. Kopp, and F. Leymann, *TOSCA: Portable Automated Deployment and Management of Cloud Applications*. New York, NY: Springer New York, 2014, pp. 527–549.
- [28] M. Zimmermann, F. W. Baumann, M. Falkenthal, F. Leymann, and U. Odefey, “Automating the provisioning and integration of analytics tools with data resources in industrial environments using opentosca,” in *Proceedings of the 2017 IEEE 21<sup>st</sup> International Enterprise Distributed Object Computing Conference Workshops and Demonstrations (EDOCW 2017)*. IEEE Computer Society, Oct. 2017, pp. 3–7.
- [29] Y. Wang, P. N. Ajoku, J. C. Brustoloni, and B. O. Nnaji, “Intellectual property protection in collaborative design through lean information modeling and sharing,” *Journal of computing and information science in engineering*, vol. 6, no. 2, 2006, pp. 149–159.
- [30] S. Li, T. Zhang, J. Gao, and Y. Park, “A sticky policy framework for big data security,” in *2015 IEEE First International Conference on Big Data Computing Service and Applications (BigDataService)*. IEEE, 2015, pp. 130–137.
- [31] T. Waizenegger et al., “Policy4TOSCA: A Policy-Aware Cloud Service Provisioning Approach to Enable Secure Cloud Computing,” in *On the Move to Meaningful Internet Systems: OTM 2013 Conferences*. Springer, Sep. 2013, pp. 360–376.
- [32] U. Breitenbücher, T. Binz, O. Kopp, F. Leymann, and M. Wieland, “Policy-aware provisioning of cloud applications,” in *Proceedings of the 7<sup>th</sup> International Conference on Emerging Security Information, Systems and Technologies (SECURWARE)*. Xpert Publishing Services (XPS), 2013, pp. 86–95.
- [33] A. A. E. Kalam et al., “Organization based access control,” in *IEEE 4<sup>th</sup> International Workshop on Policies for Distributed Systems and Networks, 2003. Proceedings. POLICY 2003*. IEEE, 2003, pp. 120–131.
- [34] M. Li, S. Yu, K. Ren, and W. Lou, “Securing personal health records in cloud computing: Patient-centric and fine-grained data access control in multi-owner settings,” in *International conference on security and privacy in communication systems*. Springer, 2010, pp. 89–106.
- [35] F. Li, B. Luo, and P. Liu, “Secure information aggregation for smart grids using homomorphic encryption,” in *2010 First IEEE International Conference on Smart Grid Communications (SmartGridComm)*. IEEE, 2010, pp. 327–332.
- [36] S. Ruj and A. Nayak, “A decentralized security framework for data aggregation and access control in smart grids,” *IEEE transactions on smart grid*, vol. 4, no. 1, 2013, pp. 196–205.
- [37] Oracle. Understanding signing and verification. [Online]. Available: <https://docs.oracle.com/javase/tutorial/deployment/jar/intro.html> [retrieved: July, 2018]
- [38] J. Huang and M. S. Fox, “An ontology of trust: formal semantics and transitivity,” in *Proceedings of the 8<sup>th</sup> international conference on electronic commerce: The new e-commerce: innovations for conquering current barriers, obstacles and limitations to conducting successful business on the internet*. ACM, 2006, pp. 259–270.
- [39] M. Falkenthal et al., “Requirements and Enforcement Points for Policies in Industrial Data Sharing Scenarios,” in *Proceedings of the 11<sup>th</sup> Advanced Summer School on Service Oriented Computing*. IBM Research Division, 2017, pp. 28–40.
- [40] F. W. Baumann, U. Breitenbücher, M. Falkenthal, G. Grünert, and S. Hudert, “Industrial data sharing with data access policy,” in *Cooperative Design, Visualization, and Engineering*. Springer International Publishing, 2017, pp. 215–219.
- [41] M. Cooper et al. *Internet X.509 Public Key Infrastructure: Certification Path Building*. [Online]. Available: <https://tools.ietf.org/html/rfc4158> [retrieved: July, 2018]
- [42] O. Kopp, T. Binz, U. Breitenbücher, and F. Leymann, “Winery – A Modeling Tool for TOSCA-based Cloud Applications,” in *Proceedings of the 11<sup>th</sup> International Conference on Service-Oriented Computing (ICSOC 2013)*. Springer, Dec. 2013, pp. 700–704.
- [43] Winery. [Online]. Available: <https://eclipse.github.io/winery/> [retrieved: July, 2018]
- [44] U. Breitenbücher, T. Binz, O. Kopp, and F. Leymann, “Vinothek – a self-service portal for toasca,” in *Proceedings of the 6<sup>th</sup> Central-European Workshop on Services and their Composition (ZEUS 2014)*. CEUR-WS.org, Feb. 2014, Demonstration, pp. 69–72.
- [45] Prototypical implementation of the secure csar concepts. [Online]. Available: <https://github.com/OpenTOSCA/winery/releases/tag/paper%2Fvfy-secure-csar> [retrieved: July, 2018]
- [46] C. D. Cera, T. Kim, J. Han, and W. C. Regli, “Role-based viewing envelopes for information protection in collaborative modeling,” *Computer-Aided Design*, vol. 36, no. 9, 2004, pp. 873–886.
- [47] L.-C. Huang, H.-C. Chu, C.-Y. Lien, C.-H. Hsiao, and T. Kao, “Privacy preservation and information security protection for patients’ portable electronic health records,” *Computers in Biology and Medicine*, vol. 39, no. 9, 2009, pp. 743–750.
- [48] G. Miklau and D. Suciu, “Controlling access to published data using cryptography,” in *Proceedings of the 29<sup>th</sup> international conference on Very large data bases-Volume 29*. VLDB Endowment, 2003, pp. 898–909.
- [49] S. D. C. di Vimercati and S. Foresti, “Privacy of outsourced data,” in *IFIP PrimeLife International Summer School on Privacy and Identity Management for Life*. Springer, 2009, pp. 174–187.
- [50] S. D. C. di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, and P. Samarati, “Encryption policies for regulating access to outsourced data,” *ACM Transactions on Database Systems (TODS)*, vol. 35, no. 2, 2010, p. 12.
- [51] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. [Online]. Available: <http://bitcoin.org/bitcoin.pdf> [retrieved: July, 2018]