

Article, Review, Communication (Type of Paper)

## Blueprinting Approach in Support of Cloud Computing

Dinh Khoa Nguyen\*, Francesco Lelli, Mike P. Papazoglou, and Willem-Jan Van Den Heuvel

European Research Institute in Service Science (ERISS), Tilburg University, Warandelaan 2, 5037 AB, Tilburg, The Netherlands

E-Mails: [D.K.Nguyen@tilburguniversity.nl](mailto:D.K.Nguyen@tilburguniversity.nl) (D-K. N.); [F.Lelli@tilburguniversity.nl](mailto:F.Lelli@tilburguniversity.nl) (F. L.); [M.P.Papazoglou@uvt.nl](mailto:M.P.Papazoglou@uvt.nl) (M-P. P.); [W.J.A.M.vdnHeuvel@uvt.nl](mailto:W.J.A.M.vdnHeuvel@uvt.nl) (W-J. vdH.)

\* Author to whom correspondence should be addressed; Email: [D.K.Nguyen@tilburguniversity.nl](mailto:D.K.Nguyen@tilburguniversity.nl).  
Tel.: +31 (0) 13 - 466 8203.

Received: / Accepted: / Published:

---

**Abstract:** Developing and deploying Service-based Applications (SBAs) using different Software-as-a-service (SaaS), Platform-as-a-service (PaaS) and Infrastructure-as-a-service (IaaS) offerings from multiple providers in the cloud is still not possible, since current cloud service offerings are often provided as monolithic, one-size-fits-all solutions and give little or no room for customization. This limits the ability of SBA developers to configure and syndicate offerings from multiple SaaS, PaaS, and IaaS providers to address their application requirements. Furthermore, combining different independent cloud services necessitates a uniform description format that facilitates the design, customization, and composition. *Cloud Blueprinting* is a novel approach that allows SBA developers to easily design, configure and deploy virtual SBA payloads on virtual machines and resource pools on the cloud. We propose the *Blueprint* concept as a uniform abstract description for cloud service offerings that may cross different cloud computing layers, i.e. SaaS, PaaS and IaaS. To support developers with the SBA design and development in the cloud, this paper introduces a formal *Blueprint Template* for unambiguously describing a blueprint, as well as a *Blueprint Lifecycle* that guides developers through the manipulation, composition and deployment of different blueprints for an SBA. Finally, the empirical evaluation of the blueprinting approach within an EC's FP7 project is reported and an associated blueprint prototype implementation is presented.

**Keywords:** Cloud Computing, Service-Based Application, Cloud Service, Blueprint

---

## 1. Introduction

The cloud abstraction model delivers a shared pool of configurable computing resources (processors, storage, etc) that can be dynamically and automatically provisioned and released. Today's cloud capabilities are defined and provided as three levels of service offerings: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). This allows users of cloud-based services to be able to focus on what the service provides them rather than how the services are implemented or hosted. Currently, there is a substantive emphasis on the low-level technological trajectories of the cloud. For instance cloud scalability is heavily researched with efforts focusing on service horizontal scaling (i.e. adding new server replicas and load balancers to distribute the load) or vertical scaling (on-the-fly changing of the assigned resources to an already running instance) and associated issues such as load balancing [1, 2, 3]. However, we observe far less concentration into other aspects of the cloud, such as cloud application development.

There is a clear need for placing emphasis on how to develop enhanced composite service offerings at the application (or SaaS)-level and assign or reassign different virtual and physical resources dynamically and elastically. This leads to forming service syndications on demand at any level of the cloud stack that may potentially involve various SaaS/PaaS/IaaS providers by breaking up the current SaaS/PaaS/IaaS monolithic approach. This article will focus on such an approach. It will first shed light on some serious shortcomings of the cloud delivery models and management approaches, and then explain how to achieve a holistic cloud solution on the basis of the concept and techniques of cloud blueprinting.

When examining the three cloud delivery models we observe a recurrent theme. They are all constrained by the capabilities that are available by the provider at their delivery level and do not allow for easy extensibility or customization options. Better ways are necessary for cloud service consumers to orchestrate a cohesive cloud computing solution and provision cloud stack services that range across networking, computing, storage resources, and applications from diverse cloud providers. To realize this vision, the key challenges that are summarized below need to be tackled.

At the IaaS level cross-configurations of the virtual machines comprising a specific cloud service are currently not possible. Moreover, cross-machine configurations need to be inferred at deployment-time instead of remaining statically allocated, as is the norm with IaaS solutions today.

Similar problems also permeate the PaaS solution space. PaaS offerings are constrained by the capabilities that are available by the PaaS provider and do not allow easy extensibility, mash up, or customization options at the PaaS consumer (or developer) level. The underlying PaaS system and the applications built on are not portable across many different public and private clouds and cloud providers. Furthermore, many hosted middleware solutions are not integrated with any IaaS management, thereby lacking elasticity and scaling benefits.

SaaS is also predominantly tethered to proprietary application platforms where the cloud provider runs all elements of the service with the client presented with a complete application. SaaS applications cannot be customized easily at the process-level and cannot be combined with external services from other providers to offer improved functionality to the client or developer.

From the preceding discussion it is apparent that the current cloud solutions are fraught with problems:

- They introduce a monolithic SaaS/PaaS/IaaS stack architecture where a one-size-fits-all mentality prevails. They do not allow developers to mix and match functionality and services from multiple application, platform and infrastructure providers and configure it dynamically to address application needs.
- They introduce rigid service orchestration practices tied to a specific resource/infrastructure configuration for the cloud services at the application level.
- The above points hamper the (re)-configuration and customization of cloud applications on demand to reflect evolving inter-organizational collaborations.

There is clearly a need to mash up services from a variety of cloud providers to create a Service-based Application (SBA). This type of integration allows the tailoring of services to specific business needs using a mixture of SaaS, PaaS and IaaS. SBAs in the cloud have been termed “BPaaS”—Business Process as a Service—reflecting the focus on enterprise specific services. BPaaS allows creating unique end-to-end business processes that are usually syndicated with other external services (possibly provided by diverse SaaS providers). BPaaS requires effective management of cloud resources, which is the topic of the following section.

Developers at the BPaaS-level need to couple their applications in whole or part with external SaaS offerings to provide opportunities related to other areas of their clients’ business. To allow full automation and optimization of BPaaS-level actions requires a causal connection of application-level process operations to configurable supporting platforms and infrastructure services.

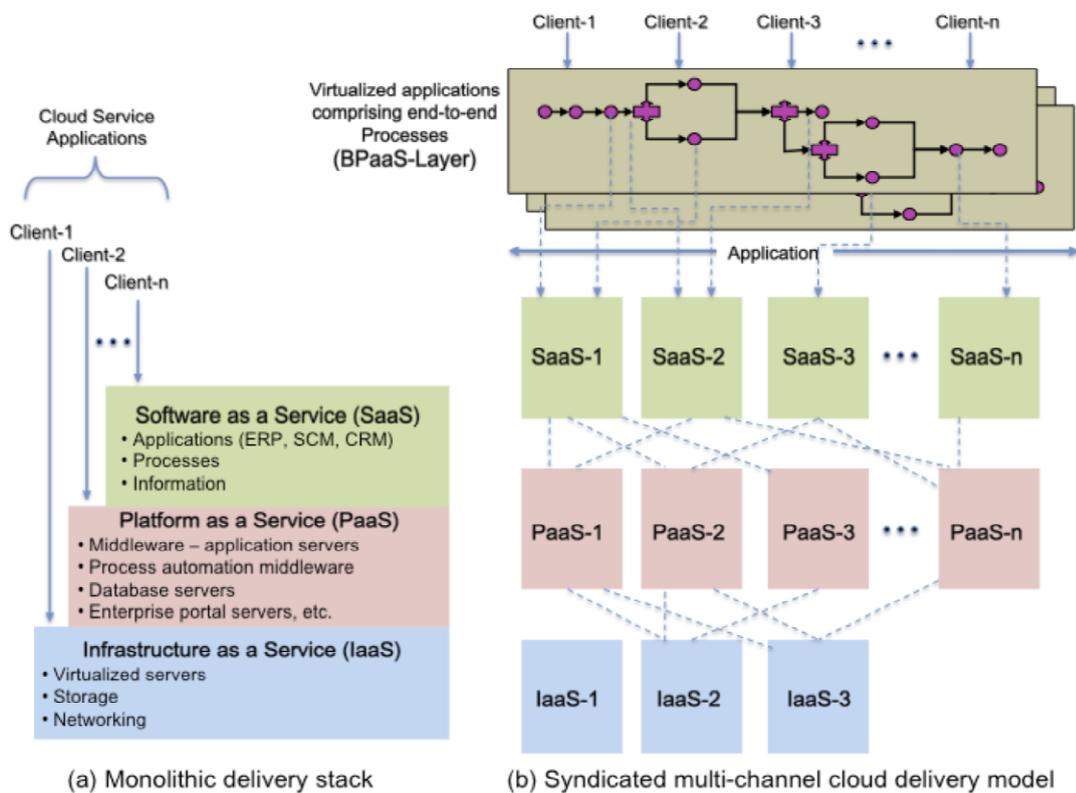


Figure 1: Different Cloud Service Delivery Models

The above approach promotes autonomous services (at all levels of the cloud stack) that adhere to the well-established software engineering principle of “separation of concerns” to minimize dependencies. This approach allows any service at any layer to be appropriately combined with a service at the same level of the cloud stack or swapped in or out without having to stop and modify other components elsewhere. At the same time, it should allow multiple (and possibly composed) resource/infrastructure or implementation options for a given service at the application-level. Such considerations lead to a *syndicated multi-channel delivery model*, which is illustrated in Figure 1 where it is contrasted with the monolithic cloud stack solutions that permeate the cloud today. Monolithic cloud solutions are shown to enforce one-way vertical deployment “channels”.

Blueprinting is a novel powerful solution allowing BPaaS applications to dynamically run on top of cloud virtualization solutions. The SaaS components of a BPaaS solution are abstracted and described in a series of templates to provide a fast, simplified method for provisioning and automating cloud services. Blueprinting seeks to simplify deployment by hiding away the complexity of deploying a BPaaS application by helping to manage all configuring of the underlying middleware and integrating with optimal PaaS and IaaS options. It achieves portability across clouds and cloud providers to leverage the benefits of elasticity and scale. Blueprinting supports a flexible top-down continuous closed-feedback loop service refinement and improvement approach. Application-level decisions regarding virtualized end-to-end services (BPaaS layer in Figure 1) are correlated with and are used to drive resource provisioning and adjust the workload and traffic to automate the dynamic configuration and deployment of application instances onto available cloud resources. The blueprinting approach promotes service virtualization by allowing an SOA-inspired approach that supports independent layering within a typical cloud stack. For example, a developer can choose to compose services from multiple SaaS providers into a coherent and integrated BPaaS, which s/he can then synthesize with platform services from one or more PaaS providers, and deploy on an IaaS cloud.

Our previous work [7] has defined a *Blueprint* as a uniform abstract description of a cloud service offering that hides all specific technical details and complexities to facilitate the SBA developers on the BPaaS layer with the selection, customization and composition of cross-layered cloud services across various syndication channels. It also proposed a structural schema called the *Blueprint Template* that can be used to describe and customize a blueprint. After releasing the first version of the template within our industrial partners in the EC’s 4caaSt FP7 project<sup>1</sup>, we have received useful comments and feedback that helped us work towards an improved version of the template in this paper. Furthermore, this paper introduces a universal *Blueprint Lifecycle* that shows how cloud-based services described by blueprints are designed, composed, and ultimately deployed on a cloud infrastructure, which will result in a running SBA in the cloud.

The rest of this paper is structured as follows: in Section 2 we introduce a cloud computing scenario that has been defined by the 4caaSt community and will be used as the running example throughout the paper; Section 3 reviews related work to show why the blueprinting approach is necessary; in Section 4 we elaborate the *Blueprint* concept, its fundamental structure and different types of blueprint; in Section 5 we propose the *Blueprint Template* to capture cross-layer cloud service offerings; Section 6 proposes a six-step lifecycle for the blueprinting approach that supports the engineering of SBAs in the

---

<sup>1</sup> EC’s 7th Framework project 4caaSt: <http://4caast.morfeo-project.org/>

cloud; in Section 7 we present our empirical evaluation approach and blueprint prototype implementation; finally, Section 8 summarizes the paper and discusses future issues.

## 2. Motivating Scenario

This section presents an enterprise cloud computing scenario developed for the EC’s 4CaaS project. Several industrial IT companies like Telefonica, Telecom Italia, Ericson, 2ndQuadrant and SAP have been involved in the definition of this scenario, which is presented as a simplified description of today industrial reality. The aim is to provide a case study that is sufficiently complex for capturing real problems and sufficiently simple for proposing and validating research solutions. The scenario promotes a uniform description for cloud service offerings to exist so that the design, configuration, and deployment of a cloud-based solution can occur. The scenario in Figure 2 is about developing a process-based application in the cloud that can receive an order for a taxi from the customer by SMS, check the current status of the taxi fleet, allocate a taxi, and finally confirm the order by SMS. It contains seven actors that collaborate in a marketplace where their cloud services can be offered and purchased. In the following, each actor is described.

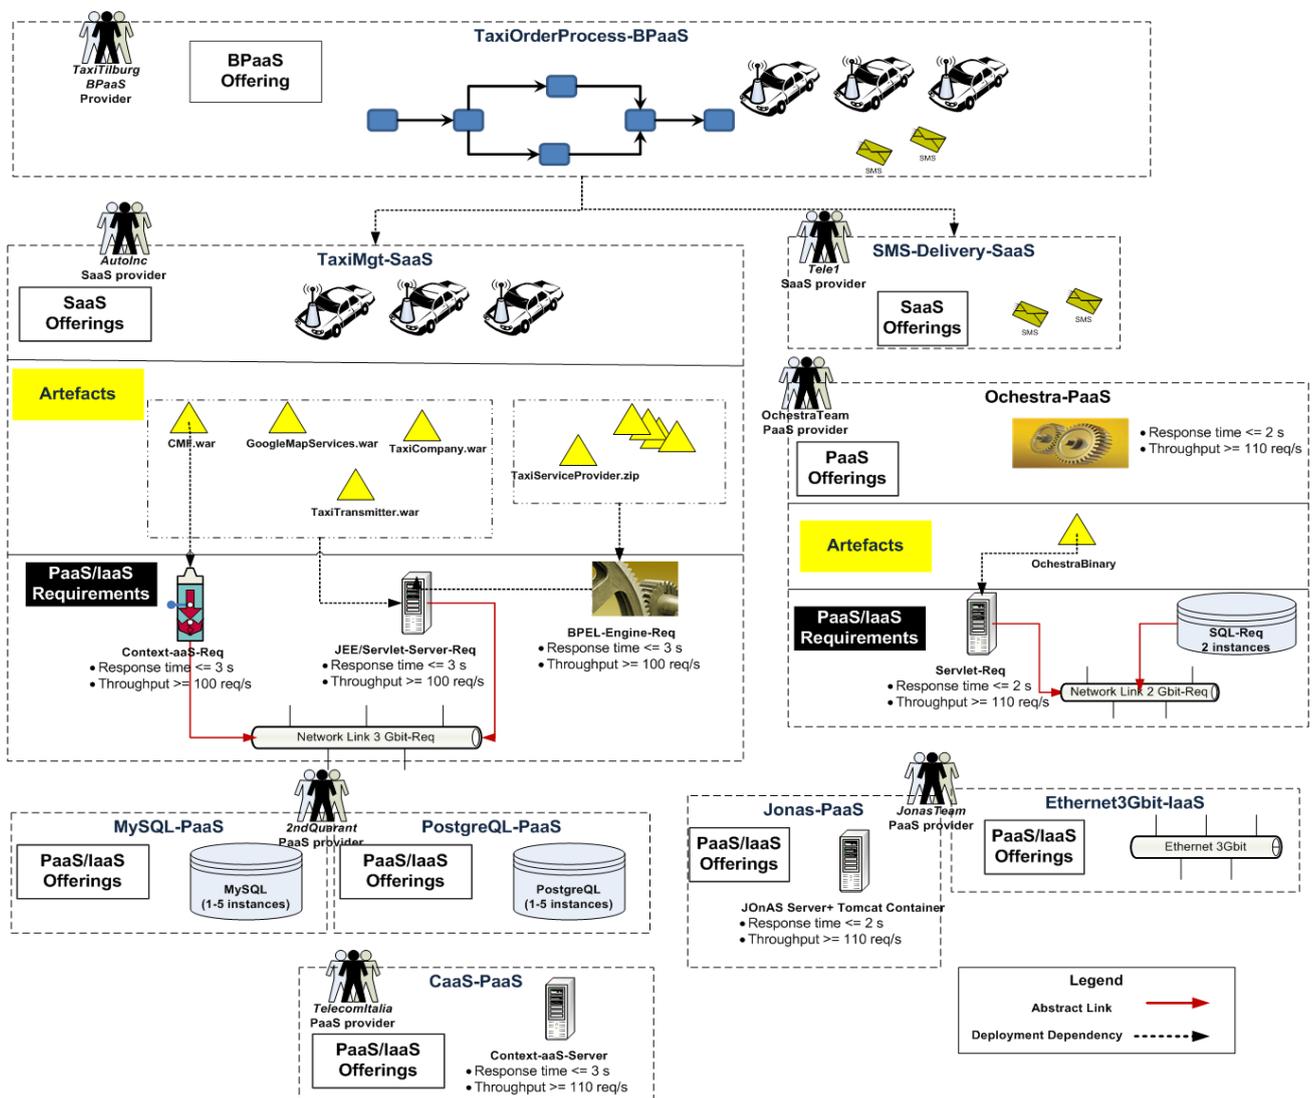


Figure 2. The 4CaaS Cloud Computing Scenario

There are three actors in the scenario that provide complete monolithic PaaS and IaaS offerings hosted on their in-house infrastructure. **2ndQuarant** is a PaaS provider specializing in storage services. It offers the relational PostgreSQL Service (*PostgreSQL-PaaS*) or the relational MySQL Service (*MySQL-PaaS*) as two alternative solutions for relational database, where each consumer can request for up to five instances of a solution. Another PaaS provider is **TelecomItalia** that provides the native Context-as-a-Service offering for context information delivery (*CaaS-PaaS*). **JonasTeam** is another PaaS provider that has two offerings in the marketplace. The first one is a JOnAS server to host Java applications including a Tomcat container to support the execution of servlets (*Jonas-PaaS*). The second one is an IaaS offering for a 3Gbit Ethernet network link (*Ethernet3Gbit-IaaS*)

**OrchestraTeam** is also a PaaS provider that offers an Orchestra BPEL Engine (*Orchestra-PaaS*). OrchestraTeam differentiates itself from other PaaS providers in a way that their offering *Orchestra-PaaS* is not complete and relies on external cloud platform and infrastructure resources for the deployment. In particular, the binary implementation file of the Orchestra engine (*OrchestraBinary*) needs a Servlet container (*Servlet-Req*) for the deployment. It is also required that the *Servlet-Req* is connected to two relational databases (*SQL-Req*) through a 2Gbit network (*NetworkLink2Gbit-Req*). To maintain the Quality of Service (QoS) promised in their offering, i.e. response time  $\leq 2s$  and throughput  $\geq 110$  request/s, OrchestraTeam requires the *Servlet-Req* to be constrained by these QoS assertions as well.

On the SaaS level, **Tele1** is a telecom company that provides a basic SMS Delivery SaaS (*SMS-Delivery-SaaS*) that is hosted on their in-house platform. In contrast, **AutoInc** is an established Medium Enterprise and has spotted a business opportunity providing fleet vehicle management in the Netherlands. They plan to deploy their business functions as a Taxi Management SaaS (*TaxiMgt-SaaS*), since this provides ubiquitous and common access for their prospective customers, e.g., taxi providers and car-hiring providers. To implement the solution, AutoInc has contracted a software consultancy who wrote the taxi management software in Java and BPEL. The software hence contains a number of war files, bpel files and other binary and configuration files (see Figure 2 for more details). The software hence requires a JEE Server (*JEE/Servlet-Server-Req*) including a Servlet v2.5 for deploying the war files for the web interface, a BPEL Engine (*BPEL-Engine-Req*) for deploying the core BPEL process files of the taxi application, and a context information provider (*Context-aaS-Req*) that is needed by the *CMF.war* file. The *BPEL-Engine-Req* is required to be deployed on the *JEE/Servlet-Server-Req*, and the *JEE/Servlet-Server-Req* is required to be connected to the *Context-aaS-Req* and to the outside through a network link with 3Gbit bandwidth (*NetworkLink3Gbit-Req*)

**TaxiTilburg** is a provider of a taxi order application that is delivered to the customers in terms of a business process (*TaxiOrderProcess-BPaaS*). Hence, their offering can be classified as a BPaaS. In order to implement the steps of their business process, TaxiTilburg completely relies on the SaaS offerings *TaxiMgt-SaaS* of AutoInc and *SMS-Delivery-SaaS* of Tele1. To maintain the predefined Quality of Service (QoS) level, TaxiTilburg insists AutoInc to specify the QoS requirements for the prospective platform's resources of the *TaxiMgt-SaaS* offering. In particular, all the required platform resources *JEE/Servlet-Server-Req*, *BPEL-Engine-Req*, *Context-aaS-Req* should ensure a throughput  $\geq 100$  req/s and a response time  $\leq 3s$ .

TaxiTilburg also has policy constraints that cannot be violated by the entire end-to-end system landscape<sup>2</sup>. These policy constraints prescribe that all the data must be stored only within the Netherlands and all the data communications must be over the Secure Socket Layer (SSL).

In summary, TaxiTilburg has designed their *TaxiOrderProcess-BPaaS* offering as an SBA, since the required *SMS-Delivery-SaaS*, *TaxiMgt-SaaS* and all the other required platform and infrastructure resources are not under their direct control. Deploying the *TaxiOrderProcess-BPaaS* on the cloud requires a syndication of the SaaS, PaaS and IaaS offerings of other providers. As we mentioned in the introduction in the previous Section 1, there is still a lack of support in building such an SBA in the cloud. Based on this case study, we will describe in the following sections the cloud blueprinting approach to select and compose different cloud service offerings to develop such an SBA.

### 3. Related Work

Much of the recent work on cloud-based application development concentrates on the infrastructure level. For example, the ability to manipulate, integrate and customize cloud service descriptions across different cloud providers has been studied in [4] that has IaaS, application and deployment orchestrators but falls short of proposing a solution for the problem at hand. The DMTF has published standards such as the Open Virtualization Format (OVF)<sup>3</sup>, to provide an open packaging and distribution format for virtual machines, and the Virtualization Management (VMAN)<sup>4</sup> specifications to address the management lifecycle of a virtual environment to help promote interoperable cloud computing services. The work in [5] is grounded on the OVF and proposes a service definition language for deploying complex Internet applications in federated IaaS clouds. These applications consist of a collection of virtual machines (VMs) with several configuration parameters (e.g., hostnames, IP addresses and other application specific details) for software components (e.g., web/application servers, database, operating system) running on the VMs. In [6], this language has been extended into a *service definition manifest* to serve as a contract between a service provider and the infrastructure provider. In this contract, architectural constraints and invariants regarding the infrastructure resource provisioning for an application service are specified and can be used for on-demand cloud infrastructure provisioning at run-time. Using the service definition manifest to specify the structure of a SaaS application, i.e. the SaaS components and their required Virtual Execution Environments (VEE), the Reservoir architecture in [8] can automatically provision the VEE instances that can run simultaneously without conflict on a federated cloud infrastructure of multiple providers. KPI monitoring mechanisms and elasticity rules in the manifest act as a contract that guarantees the required Service Level Agreement (SLA) between the SaaS provider and the Reservoir architecture.

Model-driven approaches are also employed for the purpose of automating the deployment of complex IaaS services on cloud infrastructure. For instance, the work in [9] proposes a virtual appliance model, which treats virtual images as building blocks for IaaS composite solutions. Virtual

---

<sup>2</sup> End-to-end system landscape means all of the application, platform and infrastructure systems involved in the creation of this SBA.

<sup>3</sup> DMTF: Open Virtualization Format (OVF), <http://www.dmtf.org/standards/ovf>

<sup>4</sup> DMTF: Virtual Management (VMAN), <http://dmtf.org/standards/vman>

appliances are composed into virtual solution model and deployment time requirements are then determined in a cloud-independent manner using a parameterized deployment plan. In a similar way, the approach in [10] describes a solution-based provisioning mechanism using composite appliances to automate the deployment of complex application services on a cloud infrastructure.

In conclusion, the above mentioned cloud initiatives in developing cloud-based applications, especially the ones related to the OVF, target the infrastructure level only. I.e., they allow the specification of infrastructure constraints for deploying the applications directly on (federated) data centers but fail to cover the holistic picture of a top-down development of SBAs on the cloud that can guide developers in selecting, resolving and composing SaaS and PaaS offerings.

On the cloud application level, we observe some initial efforts in proposing methodologies for application development on the cloud. The authors in [11] propose a systematic process for developing high-quality cloud SaaSs, taking into considerations the key design criteria for SaaSs and the essential commonality/variability analysis to maximize the reusability. Although this approach claims to develop cloud-based SaaSs, it does not discuss about the cloud support for the deployment environment of the SaaSs. A cloud-agnostic middleware is introduced in [12] that can sit on top of many PaaS/IaaS offerings and enable a platform-agnostic SaaS development. They provide a meta-model for describing SaaS applications and their needed cloud resources, and APIs and middleware services for the deployment. The connection between SOA and cloud computing has been studied by the Service-Oriented Cloud Computing Architecture (SOCCA) proposed in [13]. Using the SOCCA, developers can build SaaS applications following an integrated SOA framework. Cloud platform and infrastructure resources may be discovered by a Cloud Broker Layer and a Cloud Ontology Mapping Layer for deploying the SaaS components. The multi-tenancy feature of cloud computing is also supported by SOCCA where multiple instances of SaaS applications or components can be provided to multiple tenants. Although the SOCCA is a useful reference architecture for developing cloud-based SaaSs following the SOA paradigm, methodology support is lacking here including a concrete definition language for the SaaS applications and components, a composition approach for discovering and composing the cloud resources, and the ability to specify and resolve end-to-end constraints of SaaS applications that might affect the underlying cloud resources.

The Cafe application and component templates [14] are a relevant approach for cloud-based SaaS development. Cafe provides an ad-hoc composition technique for application components and cloud resources following the Service Component Architecture (SCA). However, this approach requires SaaS developers to possess intricate technical knowledge of the application architecture and the physical cloud deployment environment to select and compose the right application components and cloud resources. Furthermore, application component and cloud resource discovery in Cafe merely rely on the WS-Policy matching approach from [15] and cannot retrieve the components that satisfy end-to-end quality-of-service constraints.

In practice, an attempt to provide a template-based approach for using cloud services is available from Amazon through their AWS CloudFormation<sup>5</sup> offering. This template enables AWS developers to specify a collection of AWS cloud resources and the provisioning of these resources in an orderly

---

<sup>5</sup> Amazon Web Services: AWS CloudFormation, <http://aws.amazon.com/de/cloudformation>

and predictable fashion. Nevertheless, this template works only for AWS cloud platform and infrastructure resources and thus lacks interoperability.

In summary, existing work mostly aims to propose standards and techniques for only certain aspects and thus fails to cover the full picture of cloud computing, e.g. providing models and formats for only infrastructure resources, or describing only the functional specification. Furthermore, these standardization efforts do not aim to assist the cloud-based application developers to select, customize and compose various cross-layered cloud services across vendors according to their application requirements. We propose the *Blueprint* concept in the next section as a uniform representation to capture the comprehensive knowledge of a cloud service offering to support cloud application developers during the various development phases.

#### 4. The Blueprint Concept

Taking into account the shortcomings of existing approaches in cloud computing pointed out in Section 3, this section introduces the concept of *Blueprint* in Section 4.1 and continues by introducing different types of blueprints that can be used to capture cloud service offerings in Section 4.2.

##### 4.1 Fundamental Blueprint Structure

We define Blueprint as a uniform abstract description of a cloud service offering that abstracts away from all specific technical details and complexities to facilitate the cloud application developers with the selection, customization and composition of cross-layered cloud services across various vendors. A blueprint has the following fundamental structure, as illustrated in Figure 3:

- **Basic properties:** Some basic description properties of a blueprint including its unique ID, ownership, release date, version info, etc.
- **Offering:** The description of the cloud service offerings containing the functional capabilities, the interface specification, elasticity offering and the Quality of Service (QoS) properties.
- **Implementation Artefacts:** Description of the artifacts that implement the offering. Implementation artefacts may include the binary files, configuration files or some deployable web packages.
- **Resource Requirements:** specifying the required cloud resources for this offering including their QoS requirements. The resource requirements are needed to deploy the implementation artefacts.

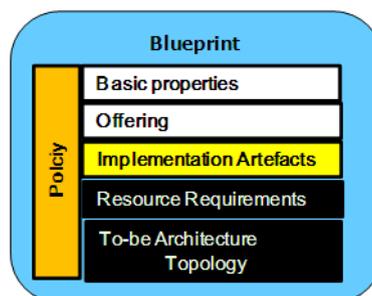


Figure 3: Fundamental Blueprint Structure

- **The To-be Architecture Topology (TAT):** specified by the application developers to capture the desired to-be architecture in terms of the interdependencies and interconnections between this blueprint and its implementation artefacts and resource requirements, or between this blueprint and other blueprints.
- **Policy:** describing the policy constraints that may not be violated by any element (the offering, implementation artefacts, resource requirements, and TAT) of a blueprint.

#### 4.2. Types of Blueprints

A blueprint can be classified based on its status. If a blueprint has already been submitted to a marketplace repository and thus can be purchased and reused by other application developers, it is called a *Source Blueprint*. In contrast, a blueprint under development is called a *Target Blueprint*. Application developers can use the blueprint template to describe their offerings in a target blueprint that will reuse other source blueprints and finally be deployed on a cloud infrastructure. For instance in the scenario in Section 2, TaxiTilburg may describe its offering through a *TaxiOrderProcess-BPaaS* target blueprint and AutoInc may describe its offering through a *TaxiMgt-SaaS* target blueprint. Other blueprints in the scenario are classified as source blueprints that are available in a marketplace repository and can be selected and reused by the application developers of TaxiTilburg or AutoInc.

Furthermore, we can also distinguish the two following types of blueprints depending on the types of their offerings, as illustrated in Figure 4:

- *Resolved blueprints:* A blueprint is resolved if its offering is already deployed and ready to be consumed. A resolved blueprint is monolithic since the offering is a predefined and pre-configured cloud stack that cannot be customized for a certain type of consumers. For instance in our scenario in Section 2, the blueprints used to describe *SMS-Delivery-SaaS*, *CaaS-PaaS*, *PostgreSQL-PaaS*, *Jonas-PaaS*, and *Ethernet3Gbit-IaaS* offerings fall under this category. Although monolithic offerings have very low flexibility for service consumers, they are currently widely adopted by the current cloud-computing providers in the market. As an example, although the CRM-related SaaSs of Salesforce<sup>6</sup> provide some customizations for the offered functionalities, they are still considered as monolithic SaaS offerings since no customizations of non-functional properties, rules or the underlying resources are allowed.
- *Unresolved blueprints:* The offering captured in this type of blueprint is incomplete, i.e. it is not yet ready to be consumed and still relies on the underlying resources for the deployment. Examples of this type of blueprint are the ones capturing the *TaxiOrderProcess-BPaaS*, *TaxiMgt-SaaS*, and *Orchestra-PaaS* offerings in the case study. In this case, the service provider needs to specify in the blueprint the resource requirements. He relies on the application developer who is responsible for the “blueprint resolution”. This task includes searching for available blueprints in the repository to fulfill the stated resource requirements. Unresolved blueprints are inherently customizable since the consumers will also be given the ability to customize the underlying configurations of the blueprints after they have been resolved. Figure 4 visualizes the unresolved blueprints on the right side (b), in which the implementation artifacts and resource requirements are now parts of the blueprint.

---

<sup>6</sup> Salesforce CRM, <http://www.salesforce.com>

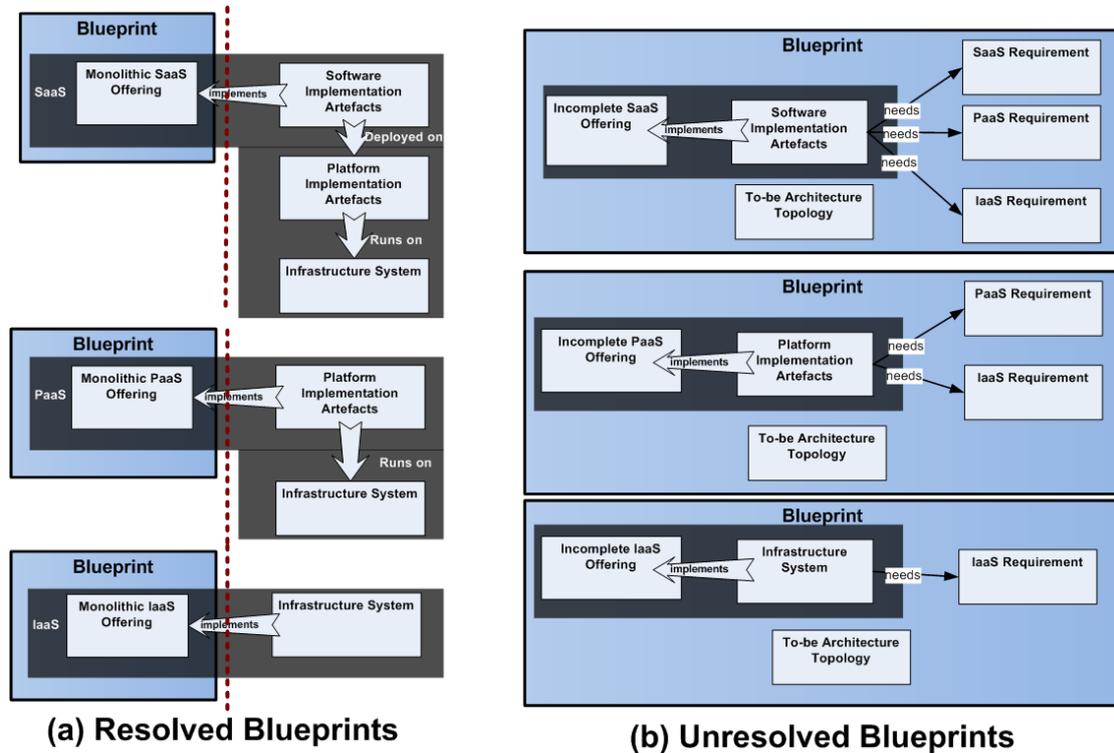


Figure 4: Resolved and Unresolved Blueprint

A target blueprint is always unresolved, as it requires other source blueprints to fulfill its resource requirements. Source blueprints can be both resolved and unresolved. In case an unresolved source blueprint is retrieved by the application developer from the marketplace for the reuse, it iteratively becomes a new target blueprint that requires the “blueprint resolution” task.

As pointed out in the introduction, our approach is to compartmentalize the monolithic cloud offering in three different layers (SaaS, PaaS, IaaS) and use the blueprints to describe each of these layered resource offerings. The advantage is that blueprints can be flexibly (re-)assembled in different configurations to form a complete SBA. The first step to enable such a flexible syndication of cross-layered blueprints is to support a formal template for describing it in an unambiguous manner. In the next Section 5 we introduce the Blueprint Template as a standard way to describe cross-layered cloud service offerings.

### 5. Blueprint Template

A blueprint has been defined in the previous Section 4 as a uniform implementation-agnostic description of a cloud service offering that abstracts away from all specific technical details and complexities. Blueprints are used to simplify the SBA development through the syndication of multi-channel composition of cross-layered cloud services across various providers. Our previous work [7] proposed a structural schema called “Blueprint Template” that can be used to describe a blueprint. After releasing the initial version of the template within the 4caast community we have received useful comments and feedback that help us work towards an improved version in this section. Figure 5 illustrates the consolidated version of the Blueprint Template. Using the template, the blueprint provider can describe (instantiate) blueprints that capture their cloud service offerings. In Figure 6, a



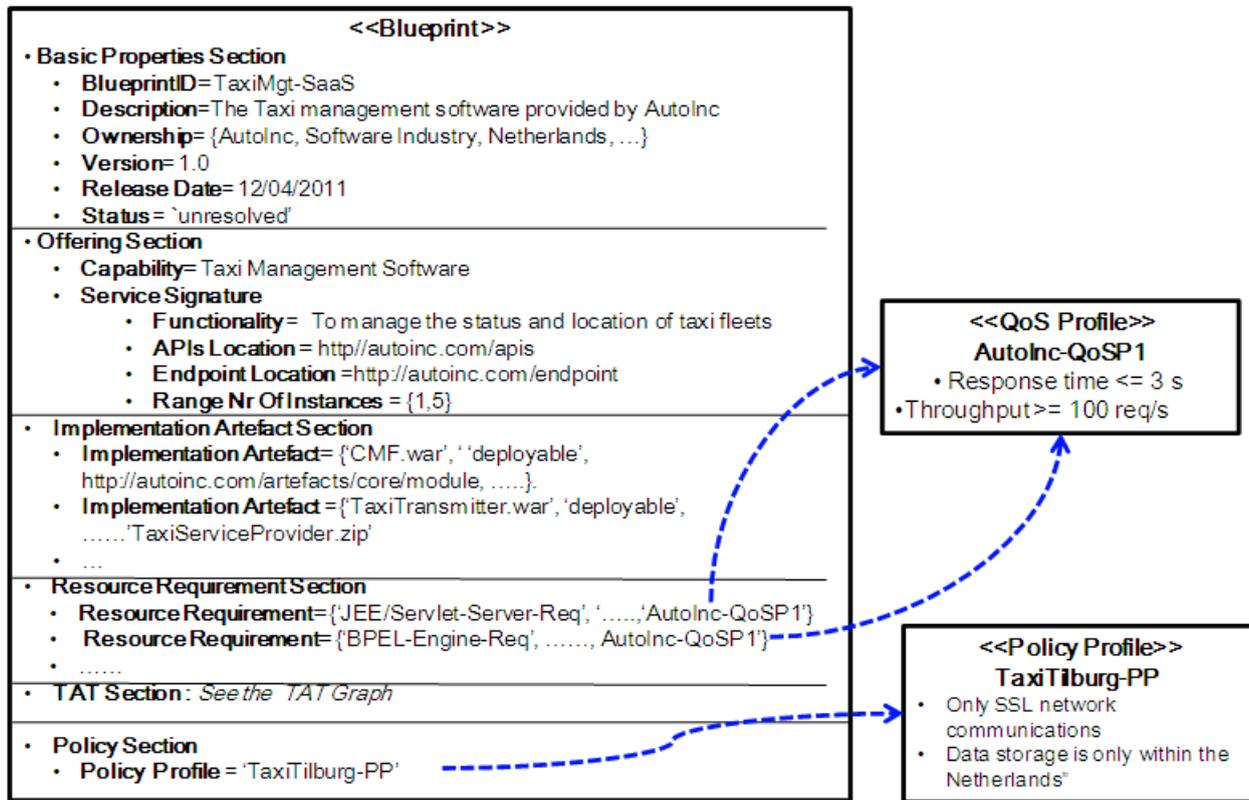


Figure 6: Example of using the Blueprint Template

### 5.2. Offering section

We assume there is only one cloud service offering in each blueprint. However, a nested offering can be referenced using the “inclusion” relationships that can be specified in another template section (see section 5.5). The functional capability of the cloud service offered in a blueprint should be described in such a way that the consumer can fully understand and query it from a blueprint repository. Initially, the capability description follows a keyword-based approach that allows only simple search and selection of blueprints. To enable more accurate search and matching, we suggest using the OVF standard for PaaS and IaaS offerings, and the service capability description template in [11] for SaaS offerings, instead of the simple keyword-based capability description.

A blueprint may have different service signatures, each of which can be described with the following information: the functionality description of the signature, the technical interfaces exposed to the blueprint consumers, the elasticity offering, and the QoS offerings. Functionality can be described following a keyword-based approach or by using external templates such as the OVF or th capability template in [11]. Technical interfaces include a location to download the APIs provided by the signature and an endpoint location for programmatic interactions with the cloud service. The APIs include not only libraries but also documentation for programming the client-side such as which protocols can be used to access the cloud service. Elasticity is specified in terms of the minimum and maximum number of instances of the cloud service (Range Nr Of Instances) provided to consumers. QoS properties of the cloud service can be specified in a number of separate profiles (QoS Profile) using an add-on templates or external languages, e.g. WS-Policy, SLAng, etc. The blueprint template allows the specification of UUID pointers referencing these separate profiles.

### 5.3. Implementation Artefacts section

This section is important for the application developer who is responsible for the deployment and provisioning of a blueprint. It encompasses the following information: an artefact id for uniquely identifying an artifact, an artefact name, an artifact type indicating whether this artefact is a software binary, a composition script, a database startup file or some other kinds of configuration files, an artifact location for downloading, and some artefact dependencies pointing to the other artefacts that have to be executed before executing this one.

Examples of implementation artefacts can be found in Figure 6. *TaxiTransmitter.war* is the ID of a deployable war file that implements parts of the Taxi application SaaS. However, it can work only after another artefact *TaxiServiceProvider.zip*, which is a zip file containing the BPEL scripts for the taxi application, has been deployed properly.

### 5.4 Resource Requirements section

A list of cloud resource requirements needed for deploying and provisioning a blueprint is specified in this section. This section supports the application developers in searching for cloud service offerings in a blueprint repository. Each resource requirement is specified with a resource ID, the required functionality, the required Range Number of Instances, the required API (if needed) and a set of references pointing to QoS Profiles that contain the QoS assertions for this resource requirement. Similarly to the offering section, the required functionality can be described using a simple keyword-based text or, in a more expressive way, using the external add-on templates like OVF or the one in [11]. The associated QoS can be described using an existing external language like WS-Policy or SLang.

As an example of resource requirements in Figure 6, the TaxiMgt-SaaS blueprint needs an instance of a JEE Application Server including a Servlet Container. This requirement has a unique id *JEE/Servlet-Server-Req* and is associated with some required QoS properties defined in the *AutoInc-Req-QoSPI* profile, which prescribes that the required resource should respond faster than 3s and provide the throughput greater than 100 requests per second.

### 5.5. To-Be Architecture Topology (TAT) section

The TAT section specifies the to-be architecture topology that is expressed through a number of relationships between blueprints, implementation artefacts and resource requirements. We use a graph data structure to capture these relationships:

- Inclusion relationships: describe the nesting relationships between blueprints. There exist also inclusion relationships between the resource requirements indicating that a required cloud resource should be provided together with another cloud resource.
- Link relationships: describe the abstract links between two resource requirements or between a blueprint and a resource requirement. The abstract links help application developers define the desired to-be architecture, i.e. how to connect the blueprints and the provisional resources together.
- Dependency relationships: This relationship indicates a deployment dependency between two elements, e.g. an implementation artefact needs a required resource for its deployment.

Figure 7 illustrates the TAT graph of the *TaxiOrderProcess-BPaaS* blueprint of TaxiTilburg in the scenario from Section 2. This graph indicates:

- The *TaxiOrderProcess-BPaaS* blueprint requires the *TaxiMgt-SaaS* blueprint of AutoInc and the *SMS-Delivery-SaaS* blueprint of Tele1.
- The *CMF.war*, *GoogleMapServices.war*, *TaxiCompany.war*, *TaxiTransmitter.war* implementation artefacts require a JEE application server that includes a Servlet v2.5 container (*JEE/Servlet-Server-Req*).
- Other implementation artefacts of the *TaxiMgt-SaaS* blueprint require a BPEL Engine (*BPEL-Engine-Req*), which is supposed to be deployed on the *JEE/Servlet-Server-Req*
- The required JEE application server (*JEE/Servlet-Server-Req*) is connected to the required Context-as-a-Service and to the outside by a network link 3Gbit (*NetworkLink3Gbit-Req*).

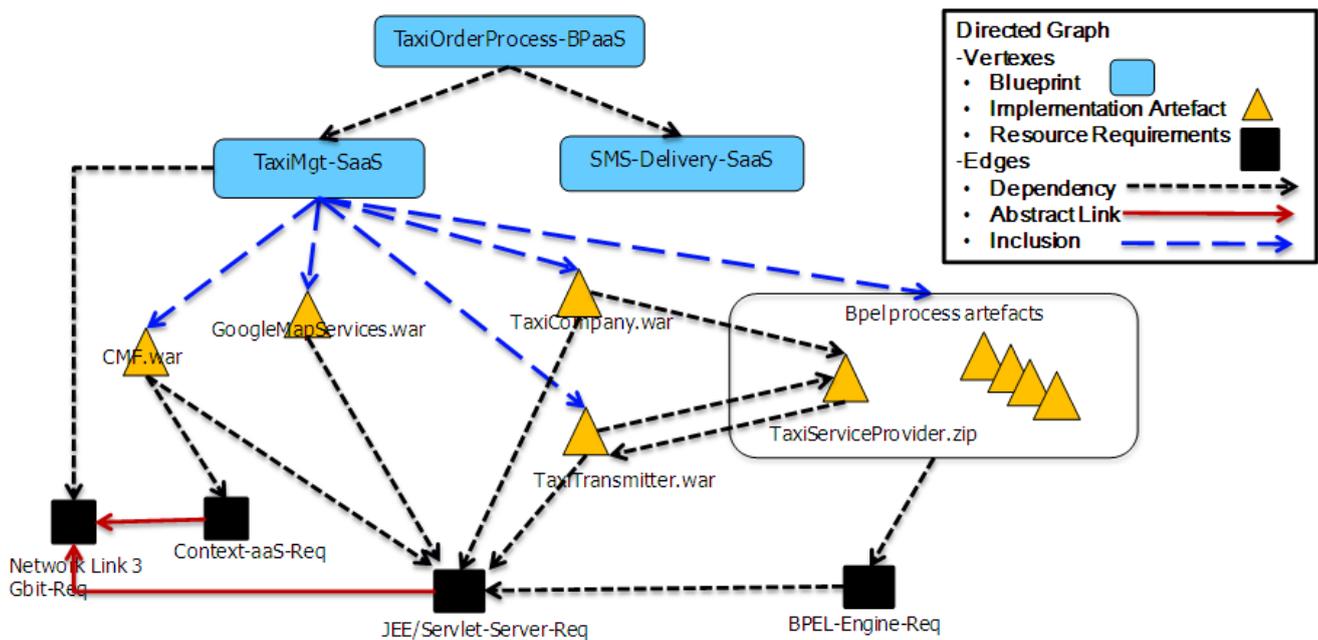


Figure 7: A sample To-be Architecture Topology (TAT) graph

In summary, the TAT is an essential part of the Blueprint Template that specifies the to-be architecture topology of a blueprint. A TAT in Figure 7 is “unresolved” as it still contains a number of resource requirements. In the further development phases of the blueprint, the resource requirements in the TAT need to be resolved by searching and selecting appropriate blueprints from the blueprint repository. The TAT becomes “resolved” if all the resource requirements can be fulfilled and hence replaced by the newly selected blueprints.

### 5.6. Policy Section

Policy constraints are the end-to-end global constraints that may not be violated by the blueprint and its resource requirements. The blueprint provider can specify the policy constraints in a Policy Profile using any existing languages such as WS-Policy, RuleML, etc.

For example in Figure 6, the *TaxiMgt-SaaS* blueprint of AutoInc is constrained by two policy constraints defined by TaxiTilburg in the policy profile *TaxiTilburg-PP*. These constraints state that all network communications have to be secured through the Secure Socket Layer (SSL) and all the data have to be stored within the Netherlands. Since they are the global end-to-end constraints specified for the *TaxiOrderProcess-BPaaS*, all the prospective blueprints that will be selected to fulfill the resource requirements are constrained by them as well.

### 6. Blueprint Lifecycle Support for Engineering Cloud-based Applications

To identify how a blueprint is specified, combined and customized with other blueprints and ultimately deployed to the cloud, this section proposes a lifecycle for engineering SBAs within which blueprints are used. Identifying the activities within the lifecycle allows us to analyze each phase the blueprint goes through to understand what it is for, what the goals and benefits are and how each activity can be supported by processes, standards and automation. The blueprint lifecycle we propose is shown in Figure 8. In particular, this figure depicts how the blueprint template is used to initiate a six-step, iterative lifecycle that results in a running cloud application. Currently, our focus is put on the first three phases of the blueprint lifecycle, basically the design, resolution and customization of a new target blueprint. The rest of the lifecycle remains as future issues and will not be discussed in depth in the following subsections.

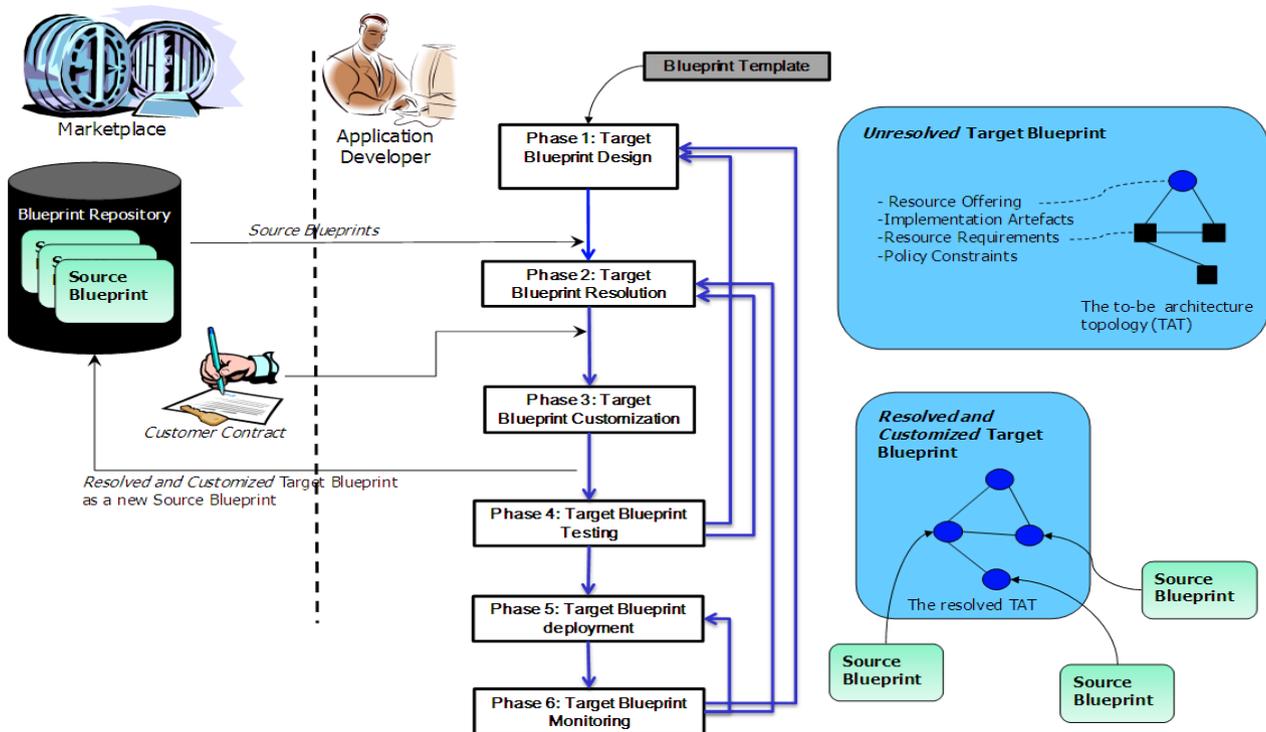


Figure 8: The Blueprint Lifecycle

#### 6.1. Phase 1: Target Blueprint Design

The blueprint lifecycle begins when a developer wishes to design a new SBA in the cloud. To do this, he uses the blueprint template, introduced in Section 5, to specify all the necessary information

about the new cloud application. The blueprint template allows the developer to describe not only the functional capabilities of the new application and the qualities it should have (e.g. minimum and maximum values for the end-to-end performance, availability and/or throughput of the complete application), but also the required resources to deploy and provision this application on the cloud. At the end of this process, when the developer has completed their specification, they have created a *Target Blueprint* that contains the application offering, the resource requirements, policy constraints, and a to-be architecture topology expressed through the relationships between the offering and requirements.

As an example of this phase, the sample *TaxiMgt-SaaS* blueprint in Section 5 is considered as a *Target Blueprint* that has been designed using our proposed blueprint template.

### 6.2. Phase 2: Target Blueprint Resolution

In order to transform the *Target Blueprint* into a description that can be deployed on cloud infrastructure, the developer must resolve the resource requirements contained in the target blueprint with concrete SaaS, PaaS or IaaS offerings of other cloud providers. These service offerings are described in *Source Blueprints*, which are specified using the same blueprint template used to develop the target blueprint described in the previous phase. The source blueprints can be found in well-known location(s) (e.g., a marketplace or repository) and are retrieved and matched to the requirements in the target blueprint using a resolution process. Matching of a source blueprint’s offering against a target blueprint’s requirement comprises of the functional matching, QoS matching and policy constraint satisfaction. We leverage the keyword-based matching approach for the functional matching between a cloud service offering and a cloud resource requirement. Since the QoS and policy properties in our blueprint may be described by the WS-Policy standard, we rely on the existing work in [24] that supports the semantic matching of WS-Policy constructs for the QoS matching and policy constraint satisfaction between the source and target blueprint.

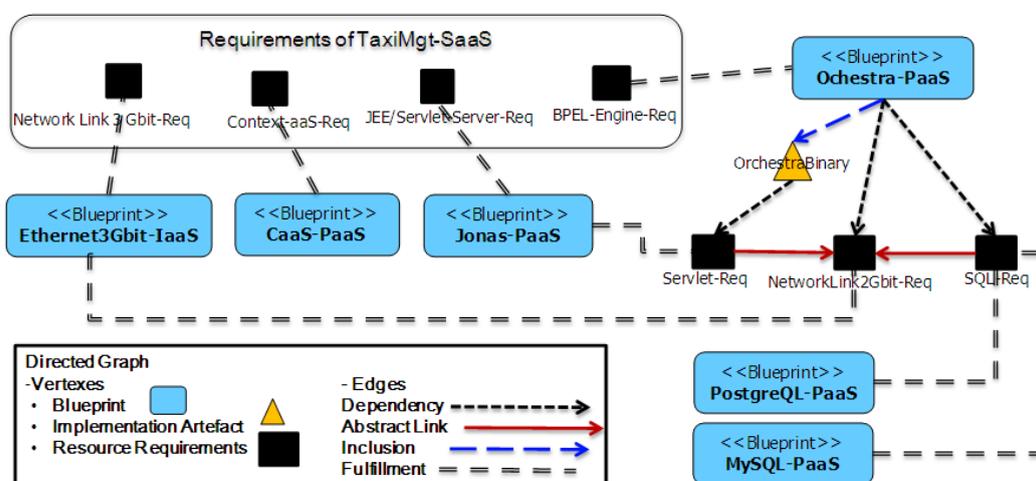


Figure 9: Resolving the sample Target Blueprint *TaxiMgt-SaaS*

Figure 9 illustrates the resolution of our sample target blueprint *TaxiMgt-SaaS*, i.e. how to fulfill its resource requirements using the offerings of other source blueprints. Note that the target blueprint resolution is an iterative process, since the source blueprints might still contain resource requirements

and iteratively becomes a new “unresolved” target blueprint. For instance in Figure 9, the source blueprint *Orchestra-PaaS* is reused to fulfill the *BPEL-Engine-Req* requirement, yet it still contains three further resource requirements that have to be fulfilled for the deployment.

The result of the resolution process is a set of alternative “resolved” target blueprints that provides options/alternatives to the developer as to how the target blueprint’s requirements can be satisfied by source blueprints (i.e., specifications of actual cloud services). For instance from Figure 9, resolving the requirement *SQL-Req* of the *Orchestra-PaaS* blueprint will result into two alternative “resolved” blueprints, one with the *PostgreSQL-PaaS* blueprint and the other with the *MySQL-PaaS* blueprint.

6.3. Phase 3: Target Blueprint Customization

The next step allows the developer to customize the target blueprint through selecting from the alternative configurations provided by the resolution process and complete the target blueprint. We assume in our case study that after resolving the *TaxiMgt-SaaS* blueprint, the application developers decide to use the PostgreSQL database instead of the MySQL database. Figure 10 shows the TAT graph of the selected “resolved” *TaxiMgt-SaaS* blueprint, in which all the resource requirements have been replaced (“resolved”) by the source blueprints.

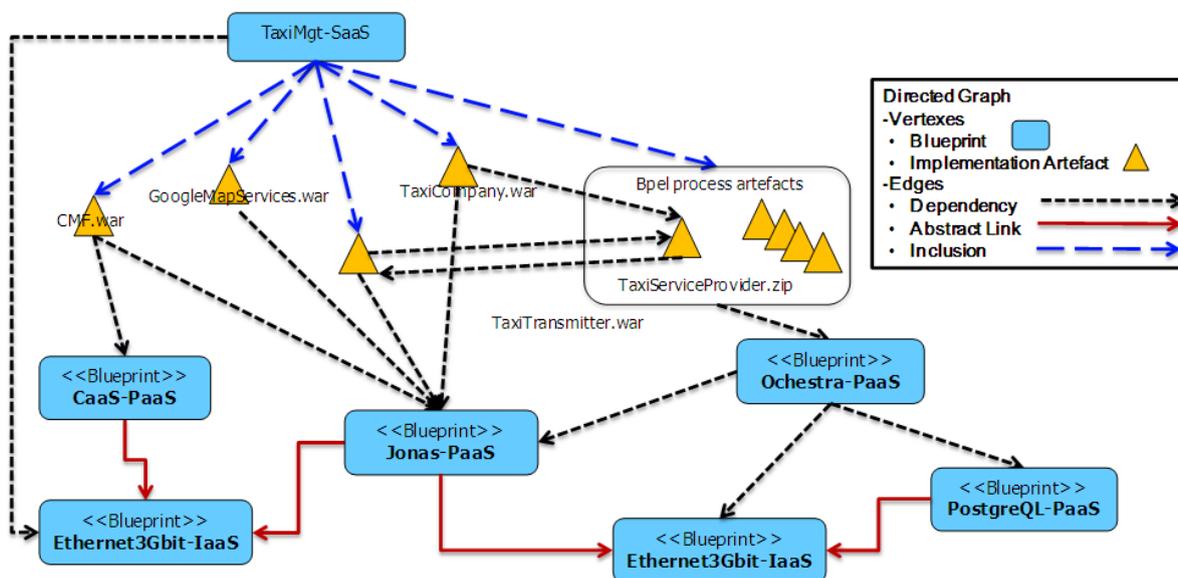


Figure 10: The TAT graph of the selected “resolved” *TaxiMgt-SaaS* Blueprint

6.4. Phase 4: Target Blueprint Checking & Testing

Once a configuration has been chosen in Phase 3, the target blueprint containing only fully-resolved blueprints is checked and tested. The purpose of this phase is to provide stakeholders with information about the behavior, properties and quality of the cloud application that has been designed, resolved and customized by the application developer from the target blueprint. If, during the testing process, the resolved and customized target blueprint fails to meet the criteria specified for the final cloud application, the developer/builder may choose one of two options: they may 1) return to Phase 1 of the lifecycle and re-design the target blueprint to create a new specification for the application, or 2) return

to Phase 2 of the lifecycle to select different source blueprints and re-configure the existing target blueprint using different offerings, perhaps from different providers.

When the application developer is satisfied that the resolved and customized target blueprint operates and behaves according to the design, meets the requirements that guided its design and development and that it can be implemented with the same characteristics the developer can choose to 1) store the resolved and customized target blueprint in the marketplace/repository so others may reuse it (possibly within a new applications) and/or 2) deploy the application using the services resolved and selected in the previous phases. However, if the application is not satisfactory and does not meet the requirements of the application developer/builder then they have two further options: 1) if the issues found during the testing process are significant, then the developer can choose to return to Phase 1 to redesign the target blueprint, possibly modifying its structure or behavior to take advantage of a particular capability; or, 2) if the problems are of a lesser magnitude, then the developer can choose to return to Phase 3 to re-select source blueprints and re-configure the original target blueprint with different cloud service offerings.

#### *6.5. Phase 5: Target Blueprint Deployment*

After the testing phase, the application is deployed onto the cloud services selected in the customization phase. The deployment is driven by a deployment plan that ensures SLAs for the correct QoS, specified and tested by the developer earlier, are agreed in advance before any of the application components are deployed. If any of the QoS cannot be agreed, the developer is notified and the customization/testing process is repeated until all QoS are satisfied. Once this step has been completed successfully, the developer, e.g., a virtual service operator or provider, deploys or reuses the SaaS, PaaS or IaaS resources required to support the application and the application is available for use.

#### *6.6. Phase 6: Target Blueprint Monitoring*

When the application has been deployed, the operational performance of the entire end-to-end process fulfilled by the application must be monitored to ensure it is performing as designed and expected, i.e., is operating within the constraints set out in the agreed SLA for the component services.

In case an SLA constraint is violated, e.g. performance has descended, a notification event should be triggered to inform the application developers to take new correcting actions. These actions could be a redeployment (which leads back to Phase 5), a new target blueprint resolution (back to Phase 2) or a completely new design of the cloud application (back to Phase 1).

In this section we have presented the lifecycle for the SBA development that incorporates iterative phases for design, testing and deployment. The benefit of documenting this process is that it demonstrates the roles, responsibilities and interactions between the developer, cloud service providers and third-parties (e.g., the marketplace or source blueprint repository) and illustrates how the blueprint template is used to specify and ultimately to deploy a running application.

### 7. Empirical Evaluation and Prototype Implementation

The empirical evaluation for the blueprinting approach aims to understand how well the proposed blueprint template can support the industry companies in describing their cloud service offerings. It started with a collaborative case study design with our industry partners in the 4caaSt project such as Telefonica, TelecomItalia, 2ndQuadrant and SAP, which resulted in the *Taxi Order Process Scenario* introduced in Section 2. Then, we implemented the proposed Blueprint Template in XSD and distributed it to the 4caaSt community to get feedback. We organized a “blueprint training” virtual workshop to explain the concept of blueprint and train our industry partners how to use the XSD blueprint template. As the cloud service offerings identified in the case study are the real industrial offerings of our partners, we requested them to use the provided XSD blueprint template to describe their blueprints in XML and submit them to our blueprint repository. Figure 11 presents a snippet of the instantiated XML document of the “unresolved” target blueprint *TaxiMgt-SaaS* from our case study. The TAT section in the *TaxiMgt-SaaS* blueprint can be visualized by the “unresolved” TAT graph on the left side of the Figure. As mentioned in Section 6.2, resolving a target blueprint means to replace the requirement IDs in its TAT template section with the IDs of the newly retrieved source blueprints. We used XQuery<sup>7</sup> for accessing and modifying data inside the XML document of the *TaxiMgt-SaaS* blueprint.

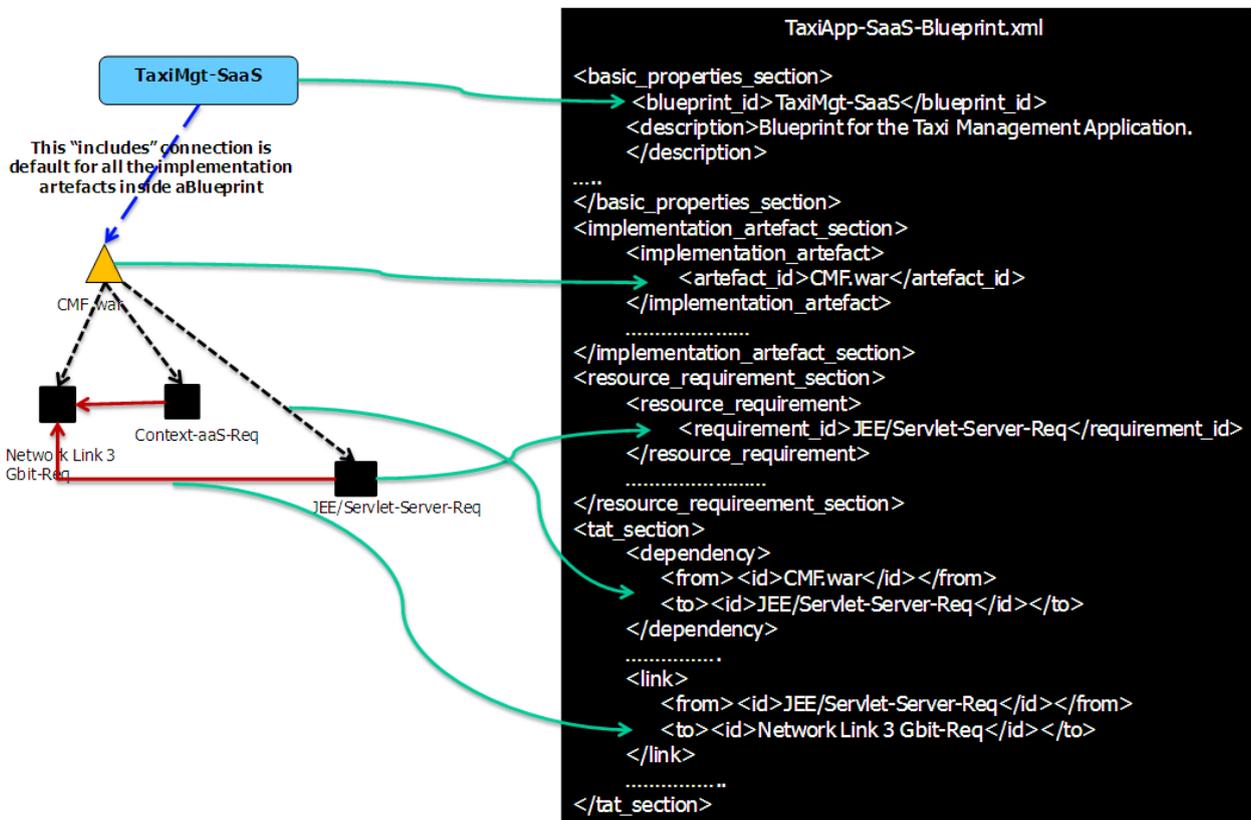


Figure 11: XML implementation of the sample *TaxiMgt-SaaS* blueprint

<sup>7</sup> <http://www.w3.org/TR/xquery/>

Our evaluation has shown that the proposed blueprint template is capable to capture all the necessary aspects of an industrial cloud service offering and simple enough to be used by our industry partners. Following our guidelines during the blueprint training workshop, a TelecomItalia representative was able to design a blueprint “on-the-fly” for his offering. We experienced just a little amount of emails exchanged for the blueprint template support, which confirms its simplicity of use.

Apart from the evaluation of the blueprint template, we have worked towards a comprehensive blueprint tool for supporting the application developers during the entire SBA development lifecycle introduced in the previous Section 6. Figure 12 depicts the prototype architecture of the blueprint tool. The Blueprint Headquarter (BlueprintHQ) server provides the SOAP API for the following three operations: *getBlueprint(BlueprintID)* to retrieve a blueprint from the repository based on its unique ID, *getBlueprints([property, value]\*)* to retrieve several blueprints based on their property values, and *resolveBlueprint(BlueprintID)* to resolve a blueprint by invoking the Blueprint Resolution Engine<sup>8</sup>. This API is consumed by the marketplace component provided by our 4caaSt partners Telefonica and SAP to productize a blueprint and advertise it in a marketplace. The API is also consumed by the BlueprintEXE server that supports a web-based user interface for visualization purpose. Figure 13 illustrates the TAT graph visualization of the sample “resolved” *TaxiMgt-SaaS* blueprint on the web interface of the BlueprintEXE. All the components of the blueprint prototype, i.e. the BlueprintEXE, BlueprintHQ, blueprint repository and blueprint resolution engine, are currently hosted on a Flexiscale virtual machine (supported by our 4caaSt partner Flexiant).

The blueprint prototype has shown the feasibility of implementing a tool for the blueprinting approach. In particular, we have merely focused on the blueprint support for the first three phases of the SBA development lifecycle described in Section 6, namely the design and reuse of blueprints for SBA development. Additional components can be integrated to the prototype including the deployment controller, monitoring component, adaptation engine, etc.

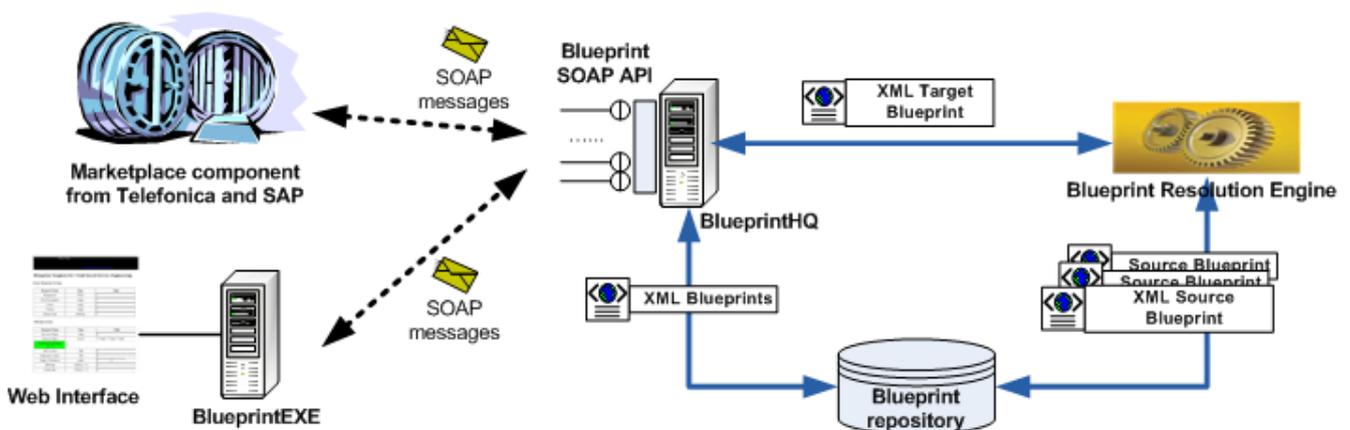


Figure 12: Blueprint Prototype Architecture

<sup>8</sup><http://blueprintexe.host-for.me/4CaaSt/blueprintHQ.htm#>

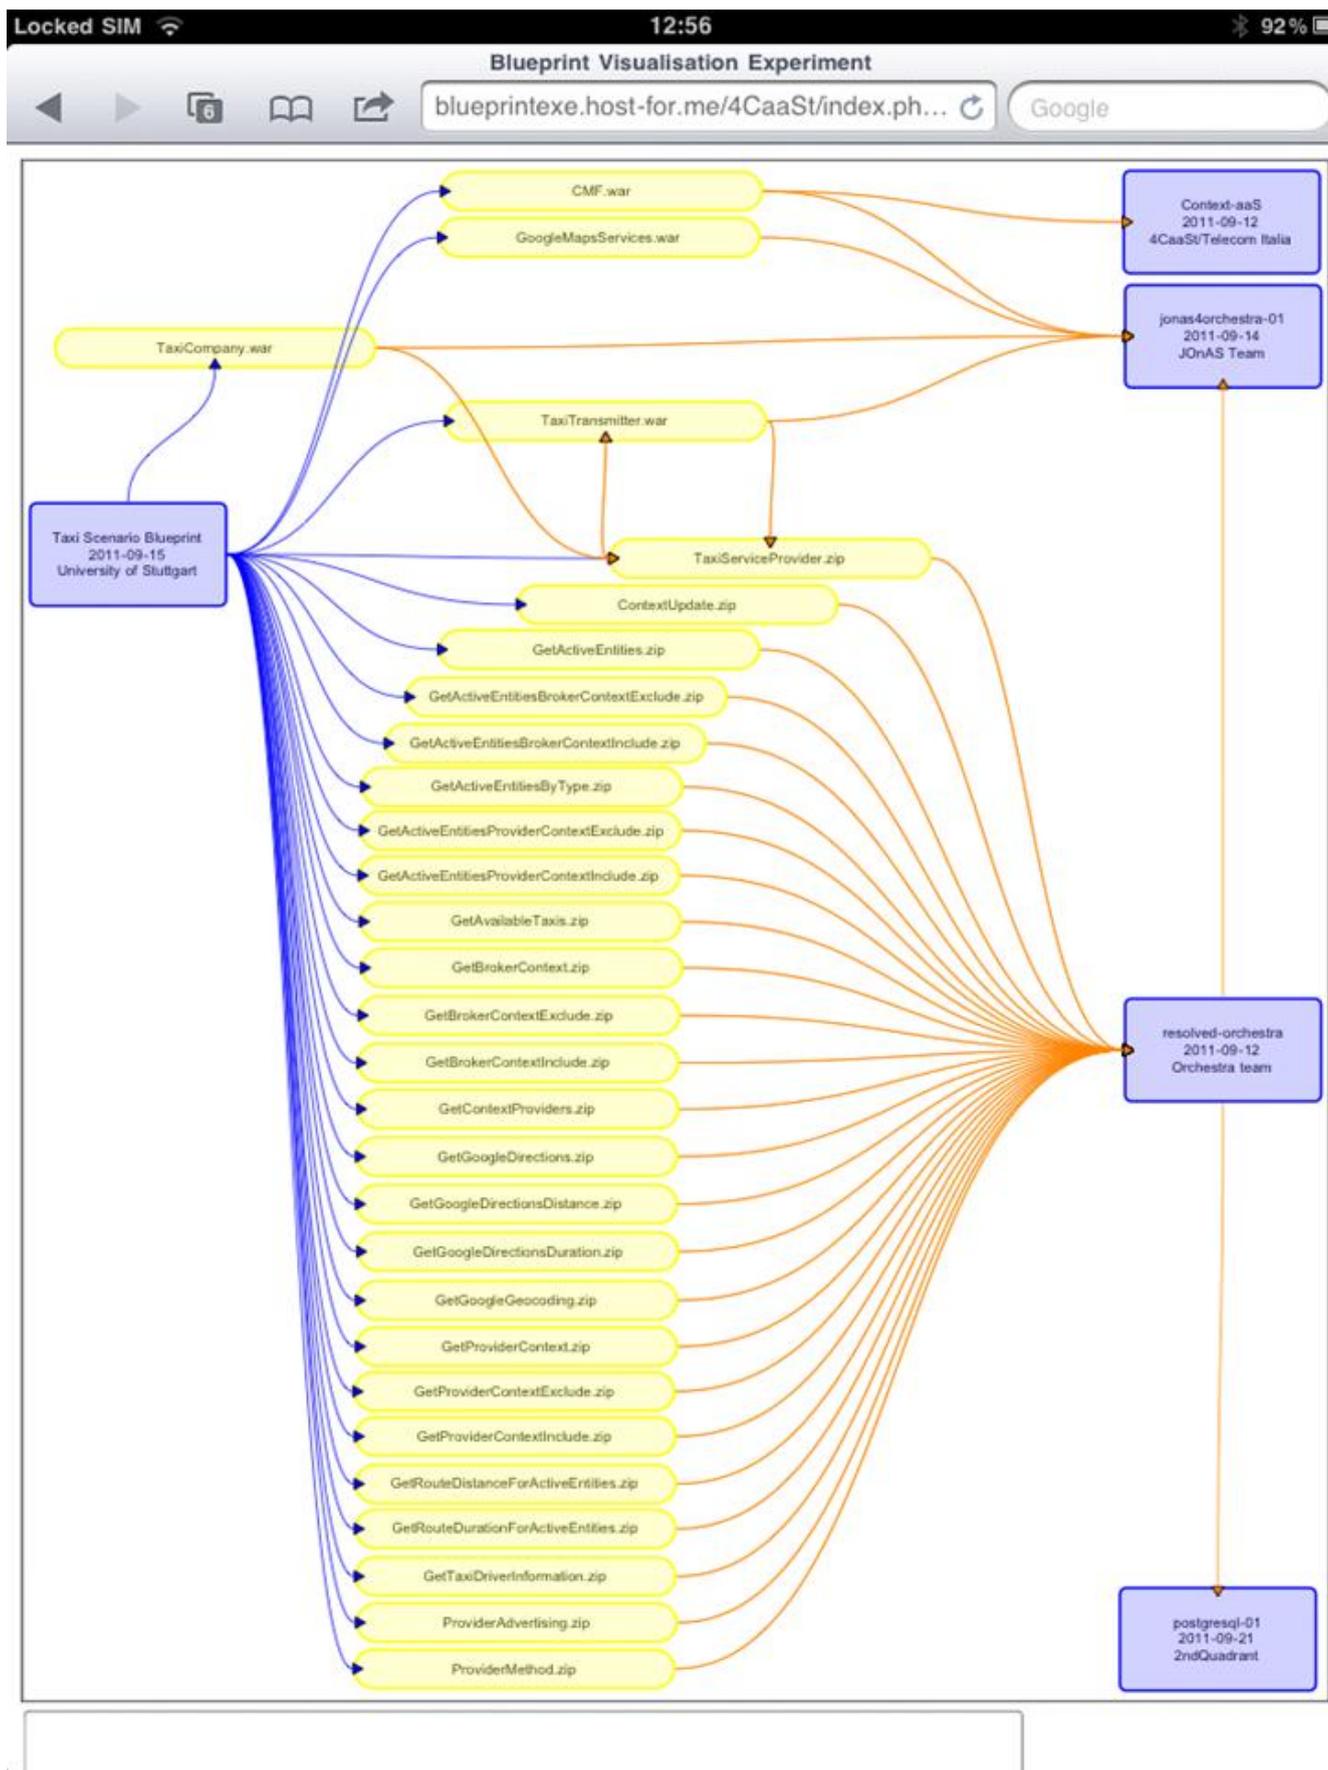


Figure 13: The web interface of our blueprint prototype showing a resolved blueprint

## 8. Conclusion and Future Work

Cloud blueprinting is a novel approach for engineering Service-based Applications (SBAs). Following this approach, developers can create sophisticated SBAs from applications, platforms and infrastructures offered by different providers in the cloud to achieve end-to-end business requirements. This paper has proposed the *Blueprint* concept as a uniform, abstract description of cross-layer cloud service offerings, a *Blueprint Template* for describing the blueprints, and a *Blueprint Lifecycle* that explains how blueprints are used during all the engineering phases of an SBA.

Blueprint has been adopted as one of the main innovative concepts within the EC's 4caaSt FP7 project. An industry cloud computing case study has been jointly defined by the 4caaSt community and has been used as a running example in this paper to demonstrate our blueprinting approach. The current blueprint XSD template and web-based blueprint prototype is integrated in a joint 4caaSt demonstration as a showcase. In the future, our vision is to continuously improve the structure of our blueprint template to capture new requirements of SBA development in the cloud, as well as develop more functionality for our blueprint prototype towards a comprehensive blueprint tool that supports the entire SBA development lifecycle.

## Acknowledgement

The research leading to this result has received funding from the Dutch Jacquard program on Software Engineering Research via contract 638.001.206 SAPIENSA; and the European Union's Seventh Framework Programme FP7/2007-2013 (4CaaSt) under grant agreement no 258862. The authors would like to thank Mathijs van der Paauw for his work on implementing the prototype.

## References and Notes

1. S. Olivier, J. Prins, "Scalable dynamic load balancing using UPC", 37th International Conference on Parallel Processing, Washington, DC, USA, IEEE Computer Society, 2008.
2. L. Rodero-Merino, et. al. "From infrastructure delivery to service management in clouds," Future Generation Computer Systems, vol. 26, October 2010.
3. H. Wu, B. Kemme, "A unified framework for load distribution and fault-tolerance of application servers," 15th International Euro-Par Conference on Parallel Processing. Berlin, Heidelberg: Springer-Verlag, 2009.
4. K. Keahey, M. Tsugawa, A. Matsunaga, and J. Fortes, "Sky computing," IEEE Internet Computing, vol. 13, no. 5, pp. 43–51, September/October 2009.
5. F. Galan, A. Sampaio, L. Rodero-Merino, I. Loy, V. Gil, and L. M. Vaquero, "Service specification in cloud environments based on extensions to open standards," in Proceedings of the Fourth International ICST Conference on COMMunication System softWARE and middlewaRE, ser. COMSWARE '09. New York, NY, USA: ACM, 2009, pp. 19:1–19:12. [Online]. Available: <http://doi.acm.org/10.1145/1621890.1621915>
6. C. Chapman, W. Emmerich, F. G. Marquez, S. Clayman, and A. Galis, "Software architecture definition for on-demand cloud provisioning," in Proceedings of the 19th ACM International

- Symposium on High Performance Distributed Computing, ser. HPDC '10. New York, NY, USA: ACM, 2010, pp. 61–72.
7. D.K. Nguyen, F. Lelli, Y. Taher, M. Parkin, M.P. Papazoglou, and W. J. van den Heuvel, “Blueprint template support for cloud-based service engineering,” in Proceedings of the Service-Wave'11, Springer LNCS (to appear), 2011.
  8. B. Rochwerger and et al., “The reservoir model and architecture for open federated cloud computing,” IBM Journal of Research and Development, vol. 53, no. 4, July 2009.
  9. A. V. Konstantinou, T. Eilam, M. Kalantar, A. A. Totok, W. Arnold, and E. Snible, “An architecture for virtual solution composition and deployment in infrastructure clouds,” in Proceedings of the 3rd international workshop on Virtualization technologies in distributed computing, ser. VTDC '09. New York, NY, USA: ACM, 2009, pp. 9–18.
  10. T. Chieu, A. Mohindra, A. Karve, and A. Segal, “Solution based deployment of complex application services on a cloud,” in Proceedings of the IEEE International Conference on Service Operations and Logistics and Informatics (SOLI), 2010.
  11. H. J. La and S. D. Kim, “A systematic process for developing high quality saas cloud services,” in Proceedings of the 1st International Conference on Cloud Computing, ser. CloudCom'09. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 278–289.
  12. E. M. Maximilien, A. Ranabahu, R. Engehausen, and L. C. Anderson, “Toward cloud-agnostic middlewares,” in Proceeding of the 24th ACM SIGPLAN conference companion on Object oriented programming systems languages and applications, ser. OOPSLA '09. New York, NY, USA: ACM, 2009, pp. 619–626.
  13. W.-T. Tsai, X. Sun, and J. Balasooriya, “Service-oriented cloud computing architecture,” 2010 Seventh International Conference on Information Technology New Generations, pp. 684–689, 2010.
  14. R. Mietzner, “A method and implementation to define and provision variable composite applications, and its usage in cloud computing,” Dissertation, Universitaet Stuttgart, Falkultaet Informatik, Elektrotechnik und Informationstechnik, Germany, August 2010.
  15. R. Mietzner, T. van Lessen, A. Wiese, M. Wieland, D. Karastoyanova, and F. Leymann, “Virtualizing services and resources with probus: The ws-policy-aware service and resource bus,” in ICWS, 2009, pp. 617–624.
  16. P. Oaks, D. Edmond, and A. ter Hofstede, “Capabilities: Describing what services can do,” in Proceedings of the First International Conference on Service-Oriented Computing, ICSOC 2003, pp 1 - 16, 2003.
  17. Kunal Verma , Rama Akkiraju , Richard Goodwin, Semantic matching of Web Service Policies, Proceedings of the Second Workshop on SDWP, 2005.