

## Smart Brix - A continuous evolution framework for Container application deployments

Johannes M Schleicher, Michael Vögler, Christian Inzinger, Schahram Dustdar

Container-based application deployments have received significant attention in recent years. Operating system virtualization based on containers as a mechanism to deploy and manage complex, large-scale software systems has become a popular mechanism for application deployment and operation. Packaging application components into self-contained artifacts has brought substantial flexibility to developers and operation teams alike. However, this flexibility comes at a price. Practitioners need to respect numerous constraints ranging from security and compliance requirements, to specific regulatory conditions. Fulfilling these requirements is especially challenging in specialized domains with large numbers of stakeholders. Moreover, the rapidly growing number of container images to be managed due to the introduction of new or updated applications and respective components, leads to significant challenges for container management and adaptation. In this paper, we introduce Smart Brix, a framework for continuous evolution of container application deployments that tackles these challenges. Smart Brix integrates and unifies concepts of continuous integration, runtime monitoring, and operational analytics. Furthermore, it allows practitioners to define generic analytics and compensation pipelines composed of self-assembling processing components to autonomously validate and verify containers to be deployed. We illustrate the feasibility of our approach by evaluating our framework using a case study from the smart city domain.

# Smart Brix – A Continuous Evolution Framework for Container Application Deployments

Johannes M. Schleicher<sup>1</sup>, Michael Vögler<sup>1</sup>, Christian Inzinger<sup>2</sup>, and Schahram Dustdar<sup>1</sup>

<sup>1</sup>Distributed Systems Group, TU Wien, Austria, {lastname}@dsg.tuwien.ac.at

<sup>2</sup>s.e.a.l. – software evolution & architecture lab, University of Zurich, Switzerland, {lastname}@ifi.uzh.ch

## ABSTRACT

Container-based application deployments have received significant attention in recent years. Operating system virtualization based on containers as a mechanism to deploy and manage complex, large-scale software systems has become a popular mechanism for application deployment and operation. Packaging application components into self-contained artifacts has brought substantial flexibility to developers and operation teams alike. However, this flexibility comes at a price. Practitioners need to respect numerous constraints ranging from security and compliance requirements, to specific regulatory conditions. Fulfilling these requirements is especially challenging in specialized domains with large numbers of stakeholders. Moreover, the rapidly growing number of container images to be managed due to the introduction of new or updated applications and respective components, leads to significant challenges for container management and adaptation. In this paper, we introduce Smart Brix, a framework for continuous evolution of container application deployments that tackles these challenges. Smart Brix integrates and unifies concepts of continuous integration, runtime monitoring, and operational analytics. Furthermore, it allows practitioners to define generic analytics and compensation pipelines composed of self-assembling processing components to autonomously validate and verify containers to be deployed. We illustrate the feasibility of our approach by evaluating our framework using a case study from the smart city domain.

Keywords: Containers, Container Evolution, Container Adaptation, DevOps, Infrastructure as Code

## 1 INTRODUCTION

In recent years, we have seen widespread uptake of operating system virtualization based on containers (Soltesz et al., 2007) as a mechanism to deploy and manage complex, large-scale software systems. Using containers, developers create self-contained images of application components along with all dependencies that are then executed in isolation on top of a container runtime (e.g., Docker<sup>1</sup>, rkt<sup>2</sup>, or Triton<sup>3</sup>). By packaging application components into self-contained artifacts, developers can ensure that the same artifact is consistently used throughout the complete software release process, from initial testing to the final production deployment. This mechanism for application deployment has become especially popular with practitioners executing projects following DevOps (Hüttermann, 2012) principles. Based on the convergence of development and operations, DevOps advocates a high degree of automation throughout the software development lifecycle (e.g., to implement continuous delivery (Humble and Farley, 2010)), along with an associated focus on deterministic creation, verification, and deployment of application artifacts using Infrastructure as Code (IaC) (Nelson-Smith, 2014) techniques, such as *Dockerfiles*<sup>4</sup> for containerized applications.

These properties allow for straightforward implementation of immutable infrastructure deployments, as advocated by IaC approaches. Application container images are usually created using a layered structure

<sup>1</sup><https://www.docker.com/>

<sup>2</sup><https://github.com/coreos/rkt>

<sup>3</sup><https://www.joyent.com/>

<sup>4</sup><https://docs.docker.com/engine/reference/builder/>

42 so that common base functionality can be reused by multiple container images. Application-specific  
43 artifacts are layered on top of a base file system so that for subsequent updates only the modified layers  
44 need to be transferred among different deployment environments. Container engine vendors such as  
45 Docker and CoreOS provide public repositories where practitioners can share and consume container  
46 images, both base images for common Linux distributions (e.g., Ubuntu, CoreOS, CentOS, or Alpine)  
47 to subsequently add custom functionality, as well as prepared application images that can be directly  
48 used in a container deployment. Once uploaded to a repository, a container image is assigned a unique,  
49 immutable identifier that can subsequently be used to deterministically deploy the exact same application  
50 artifact throughout multiple deployment stages. By deploying each application component in its own  
51 container<sup>5</sup>, practitioners can reliably execute multiple component versions on the same machine without  
52 introducing conflicts, as each component is executed in an isolated container.

53 However, since each container image must contain every runtime dependency of the packaged  
54 application component, each of these dependency sets must be maintained separately. This leads to  
55 several challenges for practitioners. Over time, the number of active container images grows due to  
56 the introduction of new applications, new application components, and updates to existing applications  
57 and their components. This growing number of container images inherently leads to a fragmentation of  
58 deployed runtime dependencies, making it difficult for operators to ensure that every deployed container  
59 continues to adhere to all relevant security, compliance, and regulatory requirements. Whenever, for  
60 instance, a severe vulnerability is found in a common runtime dependency, practitioners either have to  
61 manually determine if any active container images are affected, or initiate a costly rebuild of all active  
62 containers, irrespective of the actual occurrence of the vulnerability.

63 In this paper, we present Smart Brix, a framework for continuous evolution of container applications.  
64 Smart Brix integrates and unifies concepts of continuous integration, runtime monitoring, and operational  
65 analytics systems. Practitioners are able to define generic analytics and compensation pipelines composed  
66 of self-assembling processing components to autonomously validate and verify containers to be deployed.  
67 The framework supports both, traditional mechanisms such as integration tests, as well as custom, business-  
68 relevant processes, e.g., to implement security or compliance checks. Smart Brix not only manages the  
69 initial deployment of application containers, but is also designed to continuously monitor the complete  
70 application deployment topology to allow for timely reactions to changes (e.g., in regulatory frameworks or  
71 discovered application vulnerabilities). To enact such reactions to changes in the application environment,  
72 developers define analytics and compensation pipelines that will autonomously mitigate problems if  
73 possible, but are designed with an escalation mechanism that will eventually request human intervention  
74 if automated implementation of a change is not possible. To illustrate the feasibility of our approach we  
75 evaluate the Smart Brix framework using a case study from the smart city domain.

76 The remainder of this paper is structured as follows. In Section 2 we present a motivating scenario  
77 and relevant design goals for our framework. We present the Smart Brix framework in Section 3, along  
78 with a detailed discussion of the framework components. In Section 4 we evaluate our approach using a  
79 case study from the smart city domain. Related work is discussed in Section 5, followed by a conclusion  
80 and outlook for further research in Section 6.


## 81 2 MOTIVATION

82 In this paper, we base our discussion on a scenario containing a multi-domain expert network as created  
83 within URBEM<sup>6</sup>, a research initiative of the city of Vienna and TU Wien. To tackle the emerging  
84 complexities that arise in the smart city domain, we introduced a novel Smart City Loop (Schleicher et al.,  
85 2015b), which is depicted in Fig. 1. This loop outlines a reactive system that enables stakeholders to make  
86 informed decisions based on the models and analyses of interdisciplinary domain experts who in turn can  
87 access the large amounts of data provided by smart cities. In URBEM, a network consists of experts in  
88 the domains of energy, mobility, mathematics, building physics, sociology, as well as urban and regional  
89 planning. URBEM aims to provide decision support for industry stakeholders to plan for the future of the  
90 city of Vienna and represents a Distributed Analytical Environment (DAE) (Schleicher et al., 2015c).

91 The experts in this scenario rely on a multitude of different models and analytical approaches to  
92 make informed decisions based on the massive amounts of data that are available about the city. In turn,

<sup>5</sup>[https://docs.docker.com/engine/articles/dockerfile\\_best-practices/](https://docs.docker.com/engine/articles/dockerfile_best-practices/)

<sup>6</sup><http://urbem.tuwien.ac.at>




**Figure 1.** Smart City Loop

93 these models rely on a plethora of different tools and environments that lead to complex requirements  
 94 in terms of providing the right runtime environment for them to operate. The used tools range from  
 95 modern systems for data analytics and stream processing like Cassandra and Spark, to proprietary tools  
 96 developed by companies and research institutes with a large variance in specific versions and requirements  
 97 to run them. Additionally, these domains have to deal with a broad range of different stakeholders and  
 98 their specific security and compliance requirements. Models sometimes need to tailor their runtime  
 99 environment to specific technology stacks to ensure compliance or to be able to access the data they need.  
 100 Managing and satisfying all these requirements is a non-trivial task and a significant factor hindering  
 101 broader adoption. Therefore, this environment offers an optimal case for the advantages that come with  
 102 the use of container-based approaches. Operations teams that need to integrate these models no longer  
 103 need to be concerned with runtime specifics. Experts simply build containers that can be deployed in the  
 104 heterogenous infrastructures of participating stakeholders.

105 However, several challenges remain. In URBEM the team of experts with their plethora of different  
 106 models created over 250 different images that serve as the foundation for running containers. The models  
 107 in these containers are fueled by data from several different stakeholders in the scenario, ranging from  
 108 research institutions in the City of Vienna to industry stakeholders in the energy and mobility domain.  
 109 Each of them mandates a very distinct set of security and compliance requirements that need to be met in  
 110 order to run them. These requirements in turn are subject to frequent changes and the containers need to  
 111 be able to evolve along with them. Additionally, even though the container approach provides isolation  
 112 from the host system it is still vital to ensure that the containers themselves are not compromised. This  
 113 calls for means to check the systems running inside the container for known vulnerabilities, an issue  
 114 that is subject to heavy and fast-paced change, again requiring according evolution. A recent study<sup>7</sup>  
 115 shows that in the case of Docker, depending on the version of the images, more than 70% of the images  
 116 show potential vulnerabilities, with over 25% of them being severe. This also begs the question of  
 117 who is responsible for checking and fixing these vulnerabilities, the operations team or the experts who  
 118 created them? Despite these security and compliance constraints, the ever-changing smart city domain  
 119 itself makes it necessary for experts to stay on top of the novel toolsets that emerge in order to handle  
 120 requirements stemming from topics like Big Data or IoT. This leads to a rapid creation and adaptation  
 121 of models and their according containers, which in turn need to be checked against these constraints again.  
 122 Last but not least, these containers need to comply to certain non-functional requirements that arise from  
 123 the specific situations they are applied in. This calls for the ability to constantly check containers against  
 124 certain runtime metrics that need to be met in order to ensure that these systems are able to deliver their  
 125 expected results within stakeholder-specific time and resource constraints.

<sup>7</sup><http://www.banyanops.com/blog/analyzing-docker-hub/>



**Figure 2.** Smart Brix Framework Overview

126 All these factors lead to a complex environment that calls for an ability to easily adapt and evolve  
 127 containers to their ever-changing requirements. Specifically, we identify the following requirements in the  
 128 context of our domain:

- 129 • The ability to check a large amount of heterogenous containers against an open set of evolving  
 130 requirements. These requirements can be vulnerabilities, compliance constraints, functional tests,  
 131 or any other metric of interest for the domain.
- 132 • The ability to mitigate issues and evolve these containers based on the the results from the previously  
 133 mentioned checks.
- 134 • An approach that is applicable in the context of operations management, while still enabling the  
 135 participation of experts both for checking as well as evolution.
- 136 • An approach that can be applied to existing deployments as well as utilized to test new ones.

### 137 3 THE SMART BRUX FRAMEWORK


138 In this section, we introduce the Smart Brix framework for continuous evolution of container-based deploy-  
 139 ments, which addresses the previously introduced requirements. We start with a framework overview,  
 140 followed by a detailed description of all framework elements, and conclude with a comprehensive  
 141 description of our proof of concept implementation including possible deployment variants.

#### 142 3.1 Framework Rationales

143 The Smart Brix framework follows the microservice (Newman, 2015) architecture paradigm and an  
 144 overview of the main framework components is shown in Fig. 2. The framework is logically organized  
 145 into four main facets, which group areas of responsibility. Each of these facets is composed of multiple  
 146 components where each of these components represents a microservice. The components in the *Analyzer*  
 147 and *Compensation Facet* are managed as self-assembling components<sup>8</sup>, an approach we already success-  
 148 fully applied in previous work (Schleicher et al., 2015a). Each of these components follows the *Command*  
 149 *Pattern* (Gamma et al., 1994) and consists of multiple *processors* that are able to accept multiple inputs  
 150 and produce exactly one output. This functional approach enables a clean separation of concerns and  
 151 allows us to decompose complex problems into manageable units.

152 Fig. 3 illustrates an example of auto-assembly within the *Analyzer* facet. We see a set of *processors*,  
 153 where each processor is waiting for a specific type of input and clearly specifies the output it produces.


<sup>8</sup><http://techblog.netflix.com/2014/06/building-netflix-playback-with-self.html>



**Figure 3.** Example of auto assembling processors within the analyzer facet.

154 The processors use a message-oriented approach to exchange input and output data, where each output  
 155 and input is persistently available in the message queue and accessible by any processor. In this example  
 156 we perform an analysis of a custom-built Debian-based container that hosts the Apache HTTPD server.  
 157 There are two potential processors for the input *Artifact*, each of them able to handle a different container  
 158 format. Since in our example the *Artifact* is a *Docker Container*, only the *Docker Analyzer* reacts and  
 159 produces as output a *Docker Image*. In the next step there are two active processors, the *Docker Base*  
 160 *Image Analyzer* and the *Docker Package System Analyzer*, both taking *Docker Images* as input. Since the  
 161 *Docker Base Image Analyzer* cannot determine a base image for the given *Docker Image*, it produces  
 162 no output. However, the *Docker Package System Analyzer* is able to determine that the image uses a  
 163 *DPKG*-based package system and produces the according output. Now the *DPKG Package Analyzer*  
 164 reacts by taking two inputs, the original *Artifact* as well as the *DPKG* output and inspects the *Artifact*  
 165 via the *DPKG command* to produce a *Package List*. In the last step of this auto-assembly example the  
 166 *Vulnerability Analyzer* listens for a *Package List* and produces a *List of Vulnerabilities*. This enables a  
 167 straightforward auto-assembly approach, where connecting previous outputs to desired inputs leads to an  
 168 automatically assembled complex system consisting of simple manageable processors. This approach  
 169 further eliminates the necessity of complex composition and organization mechanisms, enabling dynamic  
 170 and elastic compositions of desired functionality, where processors can be added on demand at runtime.  
 171 This enables the previously mentioned creation of open and flexible analytics and compensation pipelines  
 172 based on this principle.

173 Additionally, the components in the analyzer and compensation facets follow the principle of *Confidence*  
 174 *Elasticity*, which means that a component or processor produces a result that is augmented with a  
 175 confidence value ( $c \in \mathbb{R}, 0 \leq c \leq 1$ ), with 0 representing no certainty and 1 representing absolute certainty  
 176 about the produced result. This allows for the specification of acceptable confidence intervals for the  
 177 framework, which augment the auto-assembly mechanism. The confidence intervals are provided as  
 178 optional configuration elements for the framework. In case the provided confidence thresholds are not



**Figure 4.** Confidence Adaptation Model Escalation

179 met, the framework follows an escalation model to find the next component or processor that is able to  
 180 provide results with higher confidence until it reaches the point where human interaction is necessary to  
 181 produce a satisfactory result (illustrated in Figure 4). Each processor ( $p_i$ ) from the set of active processors  
 182 ( $P_a$ ) provides a confidence value  $c_i$ . We define the overall confidence value of all active processors ( $c_a$ ) as  
 183  $c_a = \prod_{p_i \in P_a} c_i$ . The compensation stops when  $c_a$  meets the specified confidence interval of the framework  
 184 or a processor represents a human interaction which has a confidence value of ( $c_i = 1$ ).

### 185 3.2 Smart Brix Manager

186 In order to initiate a container evolution, the *Smart Brix Manager* is invoked via the *Smart Brix API* with  
 187 the following parameters: (i) a set of *Containers* to be inspected with (ii) the necessary *Credentials* to  
 188 analyze and evolve them, as well as an optional (iii) set of *Artifacts* necessary to compensate or analyze  
 189 the containers. In a first step the *Smart Brix Manager* queries the *Repository Manager* to see if there are  
 190 already known issues for the supplied containers. If any known issues are found, the *Smart Brix Manager*  
 191 creates a corresponding compensation topic via the messaging infrastructure by publishing the container  
 192 identifiers as well as the found issues. This represents an input that will subsequently be consumed by the  
 193 corresponding *Compensation Handlers* and starts the previously described auto-assembly process in the  
 194 *Compensation Facet*.

195 If no issues were found, the *Smart Brix Manager* hands off the supplied *Containers*, *Credentials* and  
 196 *Artifacts* to the *Dependency Manager* that is responsible for storing them in the *Dependency Repository*.  
 197 As a next step, the *Smart Brix Manager* creates a corresponding analyzer topic via the messaging  
 198 infrastructure and publishes the container identifiers to it. This generates an input that will be consumed  
 199 by the corresponding *Analyzers* and starts another auto-assembly process in the *Analyzer Facet*. The  
 200 *Smart Brix Manager* then listens to the created topic and waits for a response from the *Analyzer Facet*.  
 201 If any analyzer responds, the manager checks the confidence value of the provided results against the  
 202 configured confidence interval of the framework. If the results satisfy the interval it uses the *Repository*  
 203 *API* to store them in the *Analytics Repository*. If the confidence intervals are not satisfied, it waits for  
 204 a configured timeout for additional results to emerge. If this fails the framework escalates according to  
 205 the principle of *Confidence Elasticity* and marks the containers as required for human interaction. If the  
 206 confidence interval was met, the *Smart Brix Manager* initiates the previously mentioned auto-assembly  
 207 process in the *Compensation Facet*. The *Smart Brix Manager* then listens to the created topic and waits  
 208 for a response from any compensation handler. In case of a response, it checks the confidence values  
 209 by applying the same approach as for the *Analyzer Facet*, and stores them as compensations into the  
 210 *Analytics Repository*.

211 Furthermore, the *Smart Brix Manager* provides API endpoints to query the results of analytics and  
 212 compensation processes, as well as the current status via container identifiers.

### 213 3.3 Repository Manager

214 The *Repository Manager* provides a repository for storing analytics results of all analyzed containers as  
215 well as their corresponding compensations. The *Analytics Repository* itself is a distributed key value store  
216 that enables Analyzers as well as Compensation Handlers to store information without being bound to a  
217 fixed schema. In addition, this enables the previously mentioned open extensibility of our auto-assembly  
218 approach by allowing every component to choose the required storage format. Finally, the Repository  
219 Manager provides a service interface to store and retrieve analytics and compensation information as well  
220 as an interface for querying information based on container identifiers or other attributes.

### 221 3.4 Dependency Manager

222 The *Dependency Manager* handles necessary credentials and artifacts that are needed for processing  
223 containers. The Dependency Manager provides a service interface that allows the Smart Brix Manager to  
224 store artifacts and credentials associated with specific containers. Additionally, it provides a mechanism  
225 for components in the Analyzer and Compensation Facets to retrieve the necessary credentials and artifacts  
226 for the corresponding container IDs. Finally, it acts as service registry for components in the *Utility Facet*  
227 and exposes them to the Compensation and Analyzer Facet. The Dependency Manager uses a distributed  
228 key value store for its *Dependency Repository* in order to store the necessary information.

### 229 3.5 Utility Facet

230 The general role of the *Utility Facet* is to provide supporting services for Analyzers, Compensation  
231 Handlers, and Managers of the framework. Components in the Utility Facet register their offered  
232 services via the Dependency Manager. This provides an open and extensible approach that allows  
233 to incorporate novel elements in order to address changing requirements of container evolution. In  
234 our current architecture, the Utility Facet contains three components. First, a *Vulnerability Hub*, which  
235 represents a service interface that allows Analyzers as well as Compensation Handlers to check artifacts for  
236 vulnerabilities. The Vulnerability Hub can either utilize public repositories (e.g., the National Vulnerability  
237 Database<sup>9</sup>), or any other open or proprietary vulnerability repository. The second component is a  
238 *Compliance Hub* that allows to check for any compliance violations in the same way the Vulnerability Hub  
239 does. This is an important element in heterogenous multi-stakeholder environments, where compliance  
240 to all specified criteria must be ensured at all times. The last element is a *Metric Hub*, which allows to  
241 check artifacts for certain relevant metrics in order to ensure relevant Quality of Service constraints for  
242 containers.

### 243 3.6 Analyzers

244 The task of the components within the *Analyzer Facet* is to test containers for potential vulnerabilities,  
245 compliance violations or any other metrics. The facet is invoked by the *Smart Brix Manager*, which  
246 triggers an auto-assembly process for the given containers that should be analyzed. The Analyzer Facet  
247 can contain components for the most prominent container formats like Docker or Rkt, but due to the  
248 fact that we utilize the auto-assembly approach, we are able to integrate new container formats as they  
249 emerge. For analyzing a container an analyzer follows three basic steps: (i) Determine the base layer of  
250 the container in order to know how to access the package list. (ii) Determine the list of installed packages  
251 including their current version. (iii) Match the list of installed packages against a set of vulnerabilities,  
252 issues, or compliance constraints in order to determine the set of problems.

253 Every step can follow a different set of strategies to analyze a container represented as different  
254 processors, each of them with a specific confidence value. Possible processors for these steps are: (i) Base  
255 Image Processors, which try to determine the base layer of a container by matching their history against  
256 known base image IDs. (ii) Similarity Processors that try to select a base layer based on similarities in the  
257 history of the container with known containers by performing actions like collaborative filtering and text  
258 mining. (iii) Convention Processors that try to determine the base layer by trying common commands and  
259 checking their results. (iv) Human Provided Processors, which are human experts that manually analyze a  
260 container.

261 In order to access the containers and to perform analytics, the components within the Analyzer Facet  
262 interact with the Dependency Manager. The manager provides them with the necessary credentials for  
263 processing containers. Once the analyzers have processed a container, they publish the results, which are

---

<sup>9</sup><https://nvd.nist.gov/>

264 augmented with the confidence value, to the corresponding topic where the Smart Brix Manager carries  
265 on as previously described.

### 266 3.7 Compensation Handlers

267 The components in the *Compensation Facet* generate potential compensations for containers that have  
268 been previously identified by the Analyzers. Like the Analyzers, the *Compensation Handlers* are invoked  
269 by the Smart Brix Manager, which starts an auto-assembly process for the containers with problems  
270 that should be compensated. We provide components for the most prominent container formats, with  
271 the ability to extend the list as new formats emerge. The compensation handlers follow three basic  
272 steps: (i) Apply a compensation strategy for the container and the identified problem; (ii) Verify if the  
273 compensation strategy could be applied by rebuilding or restarting the container; (iii) Verify that the  
274 problems could be eliminated or reduced.

275 Again, every step can utilize a set of different processors, each of them with a specific confidence  
276 value, which represent different strategies. Possible processors are: (i) Container Processors, which try  
277 to use the base image's package manager to upgrade packages with identified vulnerabilities. (ii) Image  
278 Processors that try to build a new image without the vulnerabilities; (iii) Similarity Processor that try to  
279 compensate via applying steps from similar containers that do not show these vulnerabilities; (iv) Human  
280 Provided Processors, which are human experts that manually compensate a container.

281 The Compensation Handlers interact with the Dependency Manager in a similar way like the Analyzers  
282 to retrieve the necessary credentials to operate. As Image Processors and Similarity Processors build new  
283 images in order to compensate, they can request the necessary artifacts associated with an image to be  
284 able build them.

### 285 3.8 Implementation

286 We created a proof of concept prototype of our framework based on a set of RESTful microservices  
287 implemented in Ruby. Each component that exposes a service interface relies on the Sinatra<sup>10</sup> web  
288 framework. The *Repository Manager* and the *Dependency Manager* utilize MongoDB<sup>11</sup> as their storage  
289 backend, which enables the previously described distributed, open, and extendable key value store for  
290 their repositories. We implemented a *Vulnerability Hub* that uses a SQLite<sup>12</sup> storage backend to persist  
291 vulnerabilities in a structured format. It holds the recent data from the National Vulnerability Database<sup>13</sup>  
292 (NVD), specifically the listed Common Vulnerabilities and Exposures (CVEs). This CVE Hub allows  
293 to import the CVEs posted on NVD, stores them in its repository, and allows to search for CVEs by  
294 vulnerable software name as well as version via its Sinatra-based REST interface.

295 To enable the auto-assembly mechanism for each processor within each component in the *Analyzer*  
296 and *Compensation Facet*, we use a message-oriented middleware. Specifically, we utilize RabbitMQ's<sup>14</sup>  
297 topic and RPC concepts, by publishing each output and listening for its potential inputs on dedicated  
298 topics. We implemented a Docker Analyzer component with a *Base Image Processor* and a *Convention*  
299 *Processor*-based strategy. The Docker Analyzer first tries to determine the operating system distribution  
300 of the container by analyzing its history. Specifically, it uses the Docker API to generate the history for the  
301 container and selects the first layer's ID, which represents the base layer. It then matches this layer against  
302 a set of known layer IDs, which matches corresponding operating system distributions to determine which  
303 command to use for extracting the package list. If a match is found, it uses the corresponding commands  
304 to determine the package list. If the determined operating system is Ubuntu or Debian, it will use `dpkg`  
305 to determine the package list. If it was CentOS, `yum` is used, and if it was Alpine, `apk`. After parsing the  
306 package command output into a processable list of packages, it checks each package name and version by  
307 using the *CVE Hub* via its REST interface. When this step is finished the Analyzer publishes the list of  
308 possible vulnerabilities, including analyzed packages along with several runtime metrics. In case the base  
309 image strategy fails, the Docker Analyzer tries to determine the base layer including the corresponding  
310 operating system via a convention processor. Specifically, it test if the image contains any of the known  
311 package managers. Based on the results the analyzer determines the distribution flavor and continues as  
312 described above.

---

<sup>10</sup><http://www.sinatrarb.com/>

<sup>11</sup><https://www.mongodb.org/>

<sup>12</sup><https://www.sqlite.org/>

<sup>13</sup><https://nvd.nist.gov/>

<sup>14</sup><https://www.rabbitmq.com/>

313 We further implemented a Docker Compensation Handler with a *Container Processor* and an *Image*  
314 *Processor* based compensation strategy. The Container Processor tries to upgrade the container using the  
315 operating system distribution's package manager. After this operation succeeds, it checks if the number of  
316 vulnerabilities are reduced, by comparing the new version of packages against the *CVE Hub*. If this was  
317 the case it augments the results with a confidence value based on the percentage of fixed vulnerabilities  
318 and publishes the results. The Image Processor tries to fix the container by generating a new container  
319 manifest (e.g., Dockerfile). More precisely, it uses the Docker API to generate the image history and then  
320 derives a Dockerfile from this history. After this step, the Image Processor exchanges the first layer of  
321 the Dockerfile with the newest version of its base image. In cases where it cannot uniquely identify the  
322 correct Linux flavor, it generates multiple Dockerfiles, for example one for Ubuntu and one for Debian.  
323 It then checks the Dockerfiles' structure for potential external artifacts. Specifically, it searches for any  
324 *COPY* or *ADD* commands that are present in the Dockerfile. If this is the case, it contacts the Dependency  
325 Manager and attempts to retrieve the missing artifacts. Once this is finished the Image Processor tries to  
326 rebuild the image based on the generated Dockerfile. After this step is finished, the Image Processor again  
327 checks the new list of packages against the CVE Hub, and if it could improve the state of the image it  
328 publishes the results with the corresponding confidence value. The prototype implementation is available  
329 online<sup>15</sup>.

### 330 3.9 Deployment Modes

331 The Smart Brix Framework provides a container for each facet and therefore supports deployment on  
332 heterogeneous infrastructures. The framework enables wiring of components and aspects via setting the  
333 container's environment variables, enabling dynamic setups. We distinguish between two fundamental  
334 deployment modes, *Inspection Mode* and *Introspection Mode*.

#### 335 3.9.1 Inspection Mode

336 The *Inspection Mode* allows the framework to run in a dedicated inspection and compensation setting.  
337 In this mode the framework ideally runs exclusively without any other containers and utilizes the full  
338 potential of the host systems. This means that the Smart Brix Managers wait until they receive an explicit  
339 request to analyze and compensate an artifact.

#### 340 3.9.2 Introspection Mode

341 The *Introspection Mode* allows the framework to run in an active container setup. In this mode the  
342 framework constantly watches deployed containers via the Smart Brix Manager. The Manager can be  
343 provided with a list of containers to watch via a configuration setting. This provided list of containers  
344 is then analyzed and compensated. If no container lists are supplied, the Manager watches all running  
345 containers on the platform. In this case it initiates a check whenever new images are added, an image of a  
346 running container changes, or new vulnerabilities are listed in the CVE Hub.

## 347 4 EVALUATION

### 348 4.1 Setup

349 For our evaluation we used the following setup. We provisioned three instances in our private OpenStack  
350 cloud, each with 7.5GB of RAM and 4 virtual CPUs. Each of these instances was running Ubuntu 14.04  
351 LTS with Docker staged via docker-machine<sup>16</sup>. For our evaluation we choose the *inspection deployment*  
352 variant of our framework in order to stress-test the system without other interfering containers. We  
353 deployed one manager container representing the *Management Facet*, as well as two utility containers  
354 containing the *CVE Hub* and the *Messaging Infrastructure* on one instance. We then distributed 12  
355 analyzer containers with 12 compensation containers over the remaining two instances. Additionally,  
356 we deployed a cAdvisor<sup>17</sup> container on every instance to monitor the resource usage and performance  
357 characteristics of the running containers. Fig. 5 shows an overview of the deployed evaluation setup.


### 358 4.2 Experiments

359 Since we currently only have around 250 images in our URBEM setting, we extended the number of  
360 images to be evaluated. In order to get a representative set of heterogenous images we implemented a

<sup>15</sup><https://bitbucket.org/jomis/smartbrix/>

<sup>16</sup><https://docs.docker.com/machine/install-machine/>

<sup>17</sup><https://github.com/google/cadvisor>



**Figure 5.** Evaluation Setup of Smart Brix running in inspection mode

361 small service to crawl Docker Hub<sup>18</sup>. The Docker Hub is a public repository of Docker container images  
 362 of different flavors. These images range from base images, like Ubuntu and CentOS etc., to more complex  
 363 images like Cassandra and Apache Spark. We utilized the search function of the Hub to collect a set of  
 364 4000 images ordered by their popularity (number of pulls and number of stars), which ensures that we  
 365 focus on a set with a certain impact. We then extracted the name and the corresponding pull commands  
 366 along with the latest tag to form the URI of the image. This set of 4000 URIs represented the source for  
 367 our experiments, which was then split into 3 sets containing 250, 500, and 1000 images to be tested.

#### 368 4.2.1 Analyzer Experiments

369 We started our experiments with a focus on the Analyzer Facet of the framework. First, we started the  
 370 analyzer containers on one instance and started our tests with the 250 image set. After the run finished we  
 371 repeated it with the 500 and 1000 image set. After the tests with one instance, we repeated the experiments  
 372 with two instances where each run was repeated 3 times. During the tests we constantly monitored  
 373 cAdvisor to ensure that the instances were not fully utilized in order to ensure this would not skew results.

374 After the runs had finished we evaluated the vulnerability results. The analyzers logged the analyzed  
 375 images, their base image flavor (e.g. Ubuntu, Debian etc.), processing time to analyze the image, pull  
 376 time to get the image from the DockerHub as well as the overall runtime, number of packages, size of the  
 377 image, and number of vulnerabilities.


378 Over all our experiments the analyzers showed that around 93% of the analyzed images have vulnera-  
 379 bilities. This mainly stems from the fact that our implemented analyzers have a very high sensitivity and  
 380 check for any potentially vulnerable software with any potentially vulnerable configuration. However, this  
 381 does not necessarily mean that the specific combination of software and configuration in place shows the  
 382 detected vulnerability. If we only take a look at the images with a high severity according to their CVSS<sup>19</sup>  
 383 score, around 40% show to be affected which is conclusive with recent findings<sup>20</sup>. These results underline  
 384 the importance to implement the measures proposed by our framework. However, the focus of our work  
 385 and the aim of our experiments was not to demonstrate the accuracy of the implemented vulnerability  
 386 detection, but the overall characteristics of our framework, which we discuss in the remainder of this  
 387 section.

388 We first compared the overall runtime of our analyzers, specifically the difference for one instance vs  
 389 two instance deployments, the results are shown in Fig. 6a. Based on the results we see that our approach  
 390 can be horizontally scaled over two nodes leading to a performance improvement of around 40%. The fact  
 391 that in our current evaluation setting we were not able to halve the overall runtime using two instances  
 392 stems from several factors. On the one hand, we have a certain overhead in terms of management and  
 393 coordination including the fact that we only deployed one manager and storage asset. On the other hand,

<sup>18</sup><https://hub.docker.com/>

<sup>19</sup><https://nvd.nist.gov/cvss.cfm>

<sup>20</sup><http://www.banyanops.com/blog/analyzing-docker-hub/>



**Figure 6.** Evaluation results

394 a lot of the runtime is caused by the acquisition time, which is clearly bound by network and bandwidth.  
 395 Since our infrastructure is equipped with just one 100 Mbit uplink that is shared by all cloud resources,  
 396 this is a clear bottleneck. We also see that the majority of wall clock time is spent for acquisition and that  
 397 the actual processing time only amounts to approximately 3% of the overall runtime. The fact that the  
 398 acquisition time for the 1000 image set does not grow linearly like the runs with the 250 and 500 image set,  
 399 stems from Docker's image layer cache. In this case the overall acquisition time grows slower, because  
 400 a lot of images in the 1000 set share several layers, which, if already pulled by another analyzer in a  
 401 previous run, do not need to be pulled again, hence reducing the acquisition time. Finally, we demonstrate  
 402 that the average processing time of our framework is stable, which is shown in Fig. 6b. We further notice  
 403 a small increase in average processing time for the 250 image set, which is caused by the fact that this  
 404 set contains more images with larger package numbers compared to the overall amount of images tested,  
 405 resulting in a slightly higher average processing time.

#### 406 **4.2.2 Compensation Experiments**

407 In the the next part of our experiments we focused on the Compensation Facet of our framework. In order  
 408 to test the ability to automatically handle compensations of vulnerable images, we tested the implemented  
 409 Container Processor strategy. This strategy compensates found vulnerabilities via automatic upgrades of  
 410 existing images. It takes no human intervention, has a very high confidence, keeps all artifacts within the  
 411 images and is therefore optimal to test the auto-compensation ability of our framework. In the process  
 412 of compensation the Container Processor generates a new image with the upgraded packages. In order  
 413 to test this image for improvement we have to store it. This means that for every tested image we have  
 414 to hold the original image as well as its compensated version. Specifically, we choose to test the most  
 415 vulnerable images (images with the most vulnerable packages) out of the 1000 image set we tested that  
 416 are also the most prominent images in our URBEM scenario. This left us with 150 images, which we split  
 417 in three sets with 50, 100, and 150 images and started our compensation tests. We then repeated each  
 418 run to demonstrate repeatability and to balance our results. Since the Compensation Facet follows the  
 419 same principle as the Analyzer Facet we omitted testing it on one instance and immediately started with  
 420 two instances. After the tests finished, we compared the newly created images to the original ones and

421 checked if the number of vulnerabilities could be reduced.

422 Overall our experiments showed that from the 150 images we were able to auto-compensate 34 images  
423 by reducing the number of vulnerabilities. This illustrates that even a rather simple strategy leads to a  
424 significant improvement of around 22,6%, which makes this a very promising approach. In a next step,  
425 we compared the overall runtime of our compensation handlers for the three tested sets, and the results are  
426 shown in Fig. 6c. We again can clearly see that the major amount of time is spent for acquisition, in this  
427 case pulling the images that need to be compensated. The compensation itself only takes between 24%  
428 and 28% of the overall runtime and shows linear characteristics correlating with the number of images to  
429 be compensated. The comparatively low increase in acquisition time for the 150 image set again can be  
430 explained with the specific characteristics we see in Docker's layer handling.

431 In a next step, we compared the average processing time for each set, and the results are shown in  
432 Fig. 6d. We again notice similar characteristics as we saw with our analyzers. The average processing time  
433 as well as the median processing time are stable. The small increase for the 50 image set is explained with  
434 a larger number of images that contain more packages. This fact leads to relatively longer compensation  
435 times when upgrading them.

436 Our experiments showed that our framework is able to scale horizontally. We further could show  
437 that the majority of the runtime, both when analyzing and compensating images is caused by the image  
438 acquisition, and is bandwidth bound. Given the fact that in most application scenarios of our framework  
439 the images will not necessarily reside on Docker Hub, but instead in a local registry, this factor greatly  
440 relativizes. The processing time itself scales linearly with the number of analyzed packages, and the same  
441 was shown for the compensation approach. Furthermore, the processing time in our current evaluation  
442 setup is mostly constrained by the vulnerability checking and chosen storage system, which both are not  
443 the focus of our contribution. However, implementing different checking mechanism, caching of certain  
444 aspects in the utility facet, or scaling them out could reduce the overall processing time.

## 445 5 RELATED WORK

446 The rapid adoption of container-based execution environments for modern applications enables increased  
447 flexibility and fast-paced evolution. Next to this fast-paced evolution of containers, new containers are  
448 deployed whenever functionality has to be added, which leads to massive amounts of containers that need  
449 to be maintained. While the container provides an abstraction on top of the operating system, it is still  
450 vital that the underlying system complies to policies or regulations to avoid vulnerabilities. However,  
451 checking the plethora of available environments and adapting them accordingly, is not a trivial task.  
452 Among others (e.g., Arbaugh et al. (2000); Lowis and Accorsi (2009); Li et al. (2010); Yu et al. (2006)),  
453 Lowis and Accorsi (2011) propose a novel method for analyzing cloud-based services for certain types of  
454 vulnerabilities. Next to general models and methods for classifying and analyzing applications, several  
455 approaches emerged that allow vulnerability testing (e.g., Bau et al. (2010); Li et al. (2015a); Lebeau et al.  
456 (2013); Shahriar and Zulkernine (2009); Hummer et al. (2013)). Antunes and Vieira (2013) introduce  
457 SOA-Scanner, an extensible tool for testing service-based environments for vulnerabilities. Based on an  
458 iterative approach the tool discovers and monitors existing resources, and automatically applies specific  
459 testing approaches. In contrast to our approach, the aforementioned tools solely concentrate on testing  
460 and identifying possible security threats, but do not provide means for adapting the observed application  
461 or its environment accordingly.

462 More recently, container-based approaches are applied in the literature to ease development and  
463 operation of applications. Tosatto et al. (2015) analyze different cloud orchestration approaches based on  
464 containers, discuss ongoing research efforts as well as existing solutions. Furthermore, the authors present  
465 a broad variety of challenges and issues that emerge in this context. Wettinger et al. (2014) present an  
466 approach that facilitates container virtualization in order to provide an alternative deployment automation  
467 mechanism to convergent approaches that are based on idempotent scripts. By applying action-level  
468 compensations, implemented as fine-grained snapshots in the form of containers, the authors showed  
469 that this approach is more efficient, more robust, and easier to implement as convergent approaches.  
470 However, compared to our approach, the authors do not provide a framework for analyzing container  
471 application deployments, which based on identified issues triggers according compensation mechanisms.  
472 Gerlach et al. (2014) introduce Skyport, a container-based execution environment for multi-cloud scientific  
473 workflows. By employing Docker containers, Skyport is able to address software deployment challenges  
474 and deficiencies in resource utilization, which are inherent to existing platforms for executing scientific

475 workflows. In order to show the feasibility of their approach, the authors add Skyport as an extension to an  
476 existing platform, and were able to reduce the complexities that arise when providing a suitable execution  
477 environment for scientific workflows. In contrast to our approach the authors solely focus on introducing  
478 a flexible execution environment, but do not provide a mechanism for continuously evolving container-  
479 based deployments. Li et al. (2015b) present an approach that leverages Linux containers for achieving  
480 high availability of cloud applications. The authors present a middleware that is comprised of agents  
481 to enable high availability of Linux containers. In addition, application components are encapsulated  
482 inside containers, which makes the deployment of components transparent to the application. This allows  
483 monitoring and adapting components deployed in containers without modifying the application itself.  
484 Although this work shares similarities with our approach, the authors do not provide a framework for  
485 testing container-based deployments, which also supports semi-automatic compensation of found issues.

486 Next to scientific approaches, also several industrial platforms emerged that deal with the development  
487 and management of container-based applications, with the most prominent being Tutum<sup>21</sup> and Tectonic<sup>22</sup>.  
488 These cloud-based platforms allow building, deploying and managing dockerized applications. They are  
489 specifically built to make it easy for users to develop and operate the full spectrum of applications, reaching  
490 from single container apps, up to distributed microservices stacks. Furthermore, these platforms allow  
491 keeping applications secure and up to date, by providing easy patching mechanisms and holistic systems  
492 views. In contrast to our approach, these platforms only focus on one specific container technology, and  
493 are not extensible. IBM recently introduced the IBM Vulnerability Advisor<sup>23</sup>, a tool for discovering  
494 possible vulnerabilities and compliance policy problems in IBM containers. While IBM's approach  
495 shares similarities with our work, they are solely focusing on Docker containers that are hosted inside  
496 their own Bluemix environment and therefore do not provide a generic approach. Furthermore, their  
497 Vulnerability Advisor only provides guidance on how to improve the security of images, but does not  
498 support mechanisms to evolve containers.

## 499 6 CONCLUSION

500 The numerous benefits of container-based solutions have led to a rapid adoption of this paradigm in recent  
501 years. The ability to package application components into self-contained artifacts has brought substantial  
502 flexibility to developers and operation teams alike. However, to enable this flexibility, practitioners need  
503 to respect numerous dynamic security and compliance constraints, as well as manage the rapidly growing  
504 number of container images. In order to stay on top of this complexity it is essential to provide means  
505 to evolve these containers accordingly. In this paper we presented *Smart Brix*, a framework enabling  
506 continuous evolution of container application deployments. We described the URBEM scenario as a  
507 case study in the smart city context and provided a comprehensive description of its requirements in  
508 terms of container evolution. We introduced Smart Brix to address these requirements, described its  
509 architecture, and the proof of concept implementation. Smart Brix supports both, traditional continuous  
510 integration processes such as integration tests, as well as custom, business-relevant processes, e.g., to  
511 implement security, compliance, or other regulatory checks. Furthermore, Smart Brix not only enables  
512 the initial management of application container deployments, but is also designed to continuously  
513 monitor the complete application deployment topology and allows for timely reaction to changes (e.g.,  
514 discovered application vulnerabilities). This is achieved using analytics and compensation pipelines that  
515 will autonomously detect and mitigate problems if possible, but are also designed with an escalation  
516 mechanism that will eventually request human intervention if automated implementation of a change  
517 is not possible. We evaluated our framework using a representative case study that clearly showed that  
518 the framework is feasible and that we could provide an effective and efficient approach for container  
519 evolution.

520 As part of our ongoing and future work, we will extend the presented framework to incorporate more  
521 sophisticated checking and compensation mechanisms. We will integrate mechanisms from machine  
522 learning, specifically focusing on unsupervised learning techniques as a potential vector to advance  
523 the framework with autonomous capabilities. We also aim to integrate the Smart Brix framework with  
524 our work on IoT cloud applications (Inzinger et al., 2014; Vögler et al., 2015b,a). Furthermore, we

---

<sup>21</sup><https://www.tutum.co>

<sup>22</sup><https://tectonic.com>

<sup>23</sup><https://developer.ibm.com/bluemix/2015/07/02/vulnerability-advisor/>

525 plan to conduct a large-scale feasibility study of our framework in heterogenous container application  
526 deployments.

## 527 REFERENCES

- 528 Antunes, N. and Vieira, M. (2013). SOA-Scanner: An Integrated Tool to Detect Vulnerabilities in  
529 Service-Based Infrastructures. In *Proc. Intl. Conf. on Services Computing*, pages 280–287. IEEE.
- 530 Arbaugh, W. a., Fithen, W. L., and McHugh, J. (2000). Windows of vulnerability: A case study analysis.  
531 *Computer, IEEE*, 33(12):52–59.
- 532 Bau, J., Bursztein, E., Gupta, D., and Mitchell, J. (2010). State of the art: Automated black-box web  
533 application vulnerability testing. In *Proc. Symposium on Security and Privacy*, pages 332–345. IEEE.
- 534 Gamma, E., Helm, R., Johnson, R., and Vlissides, J. (1994). *Design patterns: elements of reusable*  
535 *object-oriented software*. Pearson Edu.
- 536 Gerlach, W., Tang, W., Keegan, K., Harrison, T., Wilke, A., Bischof, J., DSouza, M., Devoid, S.,  
537 Murphy-Olson, D., Desai, N., and Meyer, F. (2014). Skyport - Container-Based Execution Environ-  
538 nment Management for Multi-cloud Scientific Workflows. In *Proc. Intl. Workshop on Data-Intensive*  
539 *Computing in the Clouds*, pages 25–32. IEEE.
- 540 Humble, J. and Farley, D. (2010). *Continuous Delivery*. Pearson Edu.
- 541 Hummer, W., Rosenberg, F., Oliveira, F., and Eilam, T. (2013). Testing idempotence for infrastructure  
542 as code. In *Middleware 2013*, volume 8275 of *Lecture Notes in Computer Science*, pages 368–388.  
543 Springer.
- 544 Hüttermann, M. (2012). *DevOps for Developers*. Apress.
- 545 Inzinger, C., Nastic, S., Sehic, S., Vögler, M., Li, F., and Dustdar, S. (2014). MADCAT – a methodology  
546 for architecture and deployment of cloud application topologies. In *Proc. Intl. Symp. on Service-*  
547 *Oriented System Engineering*, pages 13–22. IEEE.
- 548 Lebeau, F., Legeard, B., Peureux, F., and Vernotte, A. (2013). Model-based vulnerability testing for web  
549 applications. In *Proc. Intl. Conf. on Software Testing, Verification and Validation Workshops*, pages  
550 445–452. IEEE.
- 551 Li, H.-C., Liang, P.-H., Yang, J.-M., and Chen, S.-J. (2010). Analysis on Cloud-Based Security Vulnera-  
552 bility Assessment. In *Proc. Intl. Conf. on E-Business Engineering*, pages 490–494.
- 553 Li, R., Abendroth, D., Lin, X., Guo, Y., Baek, H.-W., Eide, E., Ricci, R., and Van der Merwe, J. (2015a).  
554 Potassium: penetration testing as a service. In *Proc. Symposium on Cloud Computing*, pages 30–42.  
555 ACM.
- 556 Li, W., Kanso, A., and Gherbi, A. (2015b). Leveraging Linux Containers to Achieve High Availability for  
557 Cloud Services. In *Proc. Intl. Conf. on Cloud Engineering*, pages 76–83. IEEE.
- 558 Lowis, L. and Accorsi, R. (2009). On a Classification Approach for SOA Vulnerabilities. In *Proc. Intl.*  
559 *Computer Software and Applications Conf.*, pages 439–444.
- 560 Lowis, L. and Accorsi, R. (2011). Vulnerability Analysis in SOA-Based Business Processes. *IEEE Trans.*  
561 *Serv. Comput.*, 4(3):230–242.
- 562 Nelson-Smith, S. (2014). *Test-Driven Infrastructure with Chef*. O’Reilly Media, 2 edition.
- 563 Newman, S. (2015). *Building Microservices*. O’Reilly Media, Inc.
- 564 Schleicher, J. M., Vögler, M., Inzinger, C., and Dustdar, S. (2015a). Smart fabric-an infrastructure-  
565 agnostic artifact topology deployment framework. In *Proc. Intl. Conf. on Mobile Services*, pages  
566 320–327. IEEE.
- 567 Schleicher, J. M., Vögler, M., Inzinger, C., and Dustdar, S. (2015b). Towards the internet of cities: A  
568 research roadmap for next-generation smart cities. In *Proc. Intl. Workshop on Understanding the City*  
569 *with Urban Informatics*, pages 3–6. ACM.
- 570 Schleicher, J. M., Vögler, M., Inzinger, C., Hummer, W., and Dustdar, S. (2015c). Nomads - enabling  
571 distributed analytical service environments for the smart city domain. In *Proc. Intl. Conf. on Web*  
572 *Services*, pages 679–685. IEEE.
- 573 Shahriar, H. and Zulkernine, M. (2009). Automatic Testing of Program Security Vulnerabilities. In *Proc.*  
574 *Intl. Computer Software and Applications Conf.*, volume 2, pages 550–555. IEEE.
- 575 Soltesz, S., Pötzl, H., Fiuczynski, M. E., Bavier, A., and Peterson, L. (2007). Container-based operating  
576 system virtualization: a scalable, high-performance alternative to hypervisors. *SIGOPS Oper. Syst.*  
577 *Rev.*, 41(3):275–287.
- 578 Tosatto, A., Ruiu, P., and Attanasio, A. (2015). Container-Based Orchestration in Cloud: State of the Art

- 579 and Challenges. In *Proc. Intl. Conf. on Complex, Intelligent, and Software Intensive Systems*, pages  
580 70–75. IEEE.
- 581 Vögler, M., Schleicher, J. M., Inzinger, C., and Dustdar, S. (2015a). Diane - dynamic iot application  
582 deployment. In *Proc. Intl. Conf. on Mobile Services*, pages 298–305. IEEE.
- 583 Vögler, M., Schleicher, J. M., Inzinger, C., Nastic, S., Sehic, S., and Dustdar, S. (2015b). LEONORE –  
584 Large-scale provisioning of resource-constrained IoT deployments. In *Proc. Intl. Symp. on Service-  
585 Oriented System Engineering*, pages 78–87. IEEE.
- 586 Wettinger, J., Breitenbücher, U., and Leymann, F. (2014). Compensation-Based vs. Convergent Deploy-  
587 ment Automation for Services Operated in the Cloud. In *Service-Oriented Computing*, volume 8831 of  
588 *Lecture Notes in Computer Science*, pages 336–350. Springer.
- 589 Yu, W., Aravind, D., and Supthaweesuk, P. (2006). Software Vulnerability Analysis for Web Services  
590 Software Systems. In *Proc. Symposium on Computers and Communications*, pages 740–748. IEEE.