

# JCloudScale: Closing the Gap Between IaaS and PaaS

Rostyslav Zabolotnyi, Distributed Systems Group, Vienna University of Technology  
Philipp Leitner, Software Evolution & Architecture Lab, University of Zurich  
Waldemar Hummer, Distributed Systems Group, Vienna University of Technology  
Schahram Dustdar, Distributed Systems Group, Vienna University of Technology

The Infrastructure-as-a-Service (IaaS) model of cloud computing is a promising approach towards building elastically scaling systems. Unfortunately, building such applications today is a complex, repetitive and error-prone endeavor, as IaaS does not provide any abstraction on top of naked virtual machines. Hence, all functionality related to elasticity needs to be implemented anew for each application. In this paper, we present JLOUDSCALE, a Java-based middleware that supports building elastic applications on top of a public or private IaaS cloud. JLOUDSCALE allows to easily bring applications to the cloud, with minimal changes to the application code. We discuss the general architecture of the middleware as well as its technical features, and evaluate our system with regard to both, user acceptance (based on a user study) and performance overhead. Our results indicate that JLOUDSCALE indeed allowed many participants to build IaaS applications more efficiently, comparable to the convenience features provided by industrial Platform-as-a-Service (PaaS) solutions. However, unlike PaaS, using JLOUDSCALE does not lead to a loss of control and vendor lock-in for the developer.

Categories and Subject Descriptors: D.2.2.c [Software Engineering]: Distributed/Internet based software engineering tools and techniques; D.2.0.c [Software Engineering]: Software Engineering for Internet projects

General Terms: Languages, Experimentation, Performance

Additional Key Words and Phrases: Cloud Computing, Middleware, Programming, JCloudScale

## 1. INTRODUCTION

In recent years, the cloud computing paradigm [Buyya et al. 2009; Grossman 2009] has provoked a significant push towards more flexible provisioning of IT resources, including computing power, storage and networking capabilities. Besides economic factors (e.g., pay-as-you-go pricing), the core driver behind this cloud computing hype is the idea of elastic computing. Elastic applications are able to increase and decrease their resource usage based on current application load, for instance by adding and removing computing nodes. Optimally, elastic applications are cost and energy efficient (by virtue of operating close to optimal resource utilization levels), while still providing the expected level of application performance.

---

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 610802 (CloudWave) and no. 318201 (SIMPLI-CITY).

Author's addresses: R. Zabolotnyi, W. Hummer, S. Dustdar, Institute of Information Systems 184/1, Distributed Systems Group, Vienna University of Technology, Argentinierstrasse 8, A-1040 Wien, Austria. P. Leitner: s.e.a.l. - software evolution & architecture lab, University of Zurich, Binzmühlestrasse 14, 8050 Zurich, Switzerland.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© YYYY ACM 1539-9087/YYYY/01-ARTA \$15.00

DOI : <http://dx.doi.org/10.1145/0000000.0000000>

Elastic applications are typically built using either the IaaS (Infrastructure-as-a-Service) or the PaaS (Platform-as-a-Service) paradigm [Armbrust et al. 2010]. In IaaS, users rent virtual machines from the cloud provider, and retain full control (e.g., administrator rights). In PaaS, the level of abstraction is higher, as the cloud provider is responsible for managing virtual resources. In theory, this allows for more efficient cloud application development, as less boilerplate code (e.g., for creating and destroying virtual machines, monitoring and load balancing, or application code distribution) is required. However, practice has shown that today's PaaS offerings (e.g., Windows Azure, Google' AppEngine, or Amazon's Elastic Beanstalk) come with significant disadvantages, which render this option infeasible for many developers. These problems include: (1) strong vendor lock-in [Lawton 2008; Dillon et al. 2010], as one is typically required to program against a proprietary API; (2) limited control over the elasticity behavior or the application (e.g., developers have very little influence on when to scale up and down); (3) no root access to the virtual servers running the actual application code; and (4) little support for building applications that do not follow the basic architectural patterns assumed by the PaaS offering [Jayaram 2013] (e.g., Apache Tomcat based web applications). All in all, developers are often forced to fall back to IaaS for many use cases, despite the significant advantages that the PaaS model would promise.

In this paper, we introduce JCLLOUDSCALE, a Java-based middleware that eases the task of building elastic applications. Similar to PaaS, JCLLOUDSCALE takes over virtual machine management, application monitoring, load balancing, and code distribution. However, given that JCLLOUDSCALE is a client-side middleware instead of a complete hosting environment, developers retain full control over the behavior of their application. Furthermore, JCLLOUDSCALE supports a wide range of different applications. JCLLOUDSCALE applications run on top of any IaaS cloud, making JCLLOUDSCALE a viable solution to implement applications for private or hybrid cloud settings [Sotomayor et al. 2009; Abraham et al. 2010]. In summary, we claim that the JCLLOUDSCALE model is a promising compromise between IaaS and PaaS, combining many advantages of both worlds.

The main contributions of this paper are two-fold. Firstly, we describe the the JCLLOUDSCALE middleware in detail. This contribution is in extension of our initial work in [Leitner et al. 2012]. Secondly, we conducted a user study to evaluate JCLLOUDSCALE in comparison to both, existing IaaS (OpenStack and Amazon EC2) and PaaS (Amazon Elastic Beanstalk) systems. We address runtime performance impact of JCLLOUDSCALE, as well as development productivity and user acceptance. Our study results suggest that JCLLOUDSCALE increases developer productivity in comparison to pure IaaS solutions, comparable to Elastic Beanstalk. Unlike Elastic Beanstalk, JCLLOUDSCALE is more flexible, does not lead to vendor lock-in, and can also be used in a private or hybrid cloud environment. However, our results also show that there still are technical issues in the current JCLLOUDSCALE prototype that need to be addressed. Further, our results show that, in its current version, JCLLOUDSCALE indeed impacts performance in a small but noticeable manner. JCLLOUDSCALE is already available as open source project from GitHub.

The rest of this paper is structured as follows. In Section 2, we describe the basic JCLLOUDSCALE architecture, which we follow up with an in-depth discussion of specific elasticity-related features in Section 3. Section 4 gives an implementation overview of the middleware. This implementation forms the basis for the empirical evaluation in Section 5. Section 6 surveys related work, and, finally, Section 7 concludes the paper with an outlook on open issues.

## 2. THE CLOUDSCALE MIDDLEWARE

In the following, we introduce the main notions and features of JCLLOUDSCALE.

### 2.1. Basic Notions

JCLLOUDSCALE is a Java-based middleware for building elastic IaaS applications. The ultimate aim of JCLLOUDSCALE is to facilitate developers to implement cloud applications (in the following referred to as *target applications*) as local, multi-threaded applications, without even being aware of the cloud deployment. That is, the target application is not aware of the underlying physical distribution, and does not need to care about technicalities of elasticity, such as program code distribution, virtual machine instantiation and destruction, performance monitoring, and load balancing. This is achieved with a declarative programming model (implemented via Java annotations) combined with bytecode modification. To the developer, JCLLOUDSCALE appears as an additional library (e.g., a Maven dependency) plus a post-compilation build step. This puts JCLLOUDSCALE in stark contrast to most industrial PaaS solutions, which require applications to be built specifically for these platforms. Such PaaS applications are usually not executable outside of the targeted PaaS environment.

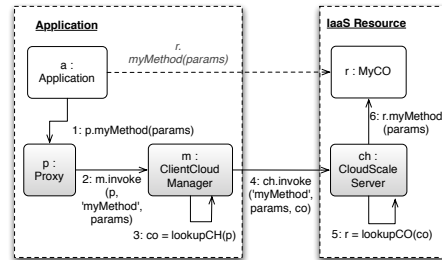


Fig. 1. Basic Interaction with Cloud Objects

The primary entities of JCLLOUDSCALE are *cloud objects* (COs). COs are object instances which execute in the cloud. COs are deployed to, and executed by, so-called *cloud hosts* (CHs). CHs are virtual machines acquired from the IaaS cloud, which run a JCLLOUDSCALE server component. They accept COs to host and execute on client request. The program code responsible for managing virtual machines, dispatching requests to virtual machines, class loading, and monitoring is injected into the target application as a post-compilation build step via bytecode modification. Optimally, COs are highly cohesive and loosely coupled to the rest of the target application, as, after cloud deployment, every further interaction with the CO constitutes a remote invocation over the network.

Fig. 1 illustrates the basic operation of JCLLOUDSCALE in an interaction diagram. The grey boxes indicate code that is injected. Hence, these steps are transparent to the application developer.

Fig. 2 shows a high-level deployment view of a JCLLOUDSCALE application. The grey box in the target application JVM again indicates injected components. Note that CHs are conceptually “thin” components, i.e., most of the actual JCLLOUDSCALE business logics is running on client side. CHs consist mainly of a small server component that accepts requests from clients, a code cache used for classloading, and sand boxes for executing COs. As JCLLOUDSCALE currently does not explicitly target multi-tenancy [Bezemer et al. 2010], these sand boxes are currently implemented in a lightweight way via custom Java classloaders. On client-side, the JCLLOUDSCALE middle-




Fig. 2. System Deployment View

ware collects and aggregates monitoring data, and maintains a list of CHs and COs. Further, the client-side middleware is responsible for scaling up and down based on user-defined policies (see Section 3.1).

```

1  @CloudObject
2  public class MyCO {
3
4      @CloudGlobal
5      private static String myCoName;
6
7      @DataSource(name = "couchdb")
8      private DataStore datastore;
9
10     @EventSink
11     private EventSink eventsink;
12
13     public MyResult myMethod(@ByValueParameter MyParameters params) {
14         ...
15     }
16 }

```

Listing 1. Declaring COs in Target Applications

## 2.2. Interacting with Cloud Objects

Application developers declare COs in their application code via simple Java annotations (see Listing 1 for a minimal example). As is the case for any object in Java, the target application can fundamentally interact with COs in two different ways: invoking CO methods, and getting or setting CO member fields. In both cases, JLOUDSCALE intercepts the operation, executes the requested operation on the CH, and returns the result (if any) back to the target application. In the meantime, the target application is blocked (more concretely, the target application remains in an “idle wait” state while it is waiting for the CH response). Fundamentally, JLOUDSCALE aims to preserve the functional semantics of the target application after bytecode modification. That is, every method call or field operation behaves functionally identical to a regular Java program.

One exception to this rule are CO-defining classes that contain static fields and methods. Operations on those are by default not intercepted by JCLLOUDSCALE, as they potentially lead to a problem that we refer to as *JVM-local updates*: if code executing on a CH (for instance a CO instance method) changes the value of a static field, only the copy in this CH's JVM will be changed. Other COs, or the target application JVM, are not aware of the change. Hence, in this case, the value of the static field is tainted, and the execution semantics of the application changes after JCLLOUDSCALE bytecode injection. To prevent this problem and preserve standard Java language semantics, static fields can be annotated with the `@CloudGlobal` annotation (see Listing 1, line 4-5). Changes to such static fields are maintained in the target application JVM, and all CH JVMs are operating on the target application JVM copy via callback. Note that this behavior is not default for performance reasons, as synchronizing static field values is expensive, and only required if JVM-local updates are possible.

### 2.3. Remote Classloading

Whenever a CH has to execute a CO method, JCLLOUDSCALE has to ensure that all necessary resources (i.e., program code and other files, for instance configuration files) are available on that CH. In order to ensure freshness of the available code and to retrieve missing files, we intercept the default class loading mechanism of Java and verify that the code available to the CH is the same as the one referenced by the client. If this is not the case, the correct version of the code is fetched dynamically from the target application. In order to improve performance, CHs additionally maintain a code cache, which is a high-speed storage of recently used code. This mechanism allows JCLLOUDSCALE to load missing or modified code efficiently and seamlessly for the application only whenever it is necessary, thus simplifying application development and maintenance. We discuss this process in more detail in [Zabolotnyi et al. 2013].

## 3. SUPPORTING CLOUD ELASTIC APPLICATIONS

So far, we have discussed how JCLLOUDSCALE transparently enables remoting in cloud applications. We now explain how JCLLOUDSCALE enables elastic applications.

### 3.1. Autonomic Elasticity via Complex Event Processing

One central advantage of JCLLOUDSCALE is that it allows for building elastic applications by mapping requests to a dynamic pool of CHs. This encompasses three related tasks: (1) performance monitoring, (2) CH provisioning and de-provisioning, and (3) CO-to-CH scheduling and CO migration. One design goal of JCLLOUDSCALE is to abstract from technicalities of these tasks, but still grant developers low-level control over the elasticity behavior.

An overview over the JCLLOUDSCALE components related to elasticity, and their interactions, is given in Fig. 3. Conceptually, our system implements the well-established autonomic computing control loop of monitoring-analysis-planning-execution [Kephart and Chess 2003] (MAPE). The base data of monitoring is provided using event messages. All components in a JCLLOUDSCALE system (COs, CHs, as well as the middleware itself) trigger a variety of predefined lifecycle and status events, indicating, for instance, that a new CO has been deployed or that the execution of a CO method has failed. Additionally, JCLLOUDSCALE makes it easy for applications to trigger custom (application-specific) events. Finally, events may also be produced by *external event sources*, such as an external monitoring framework. All these events form a consolidated stream of monitoring events in a *message queue*, by which they are forwarded into a *complex event processing (CEP) engine* [Luckham 2002] for analysis. CEP is the process of merging a large number of low-level events into high-level knowledge,

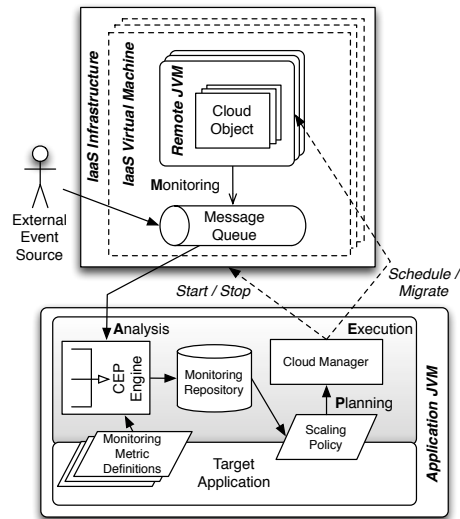


Fig. 3. Autonomic Elasticity

e.g., many atomic execution time events can be merged into meaningful performance indicators for the system in total.

Developers steer the scaling behavior by defining a *scaling policy*, which implements the planning part of this MAPE loop. This policy is invoked whenever a new CO needs to be scheduled. Possible decisions of the *scaling policy* are the provisioning of new CHs, migrating existing COs between CHs, and scheduling the new CO to a (new or existing) CH. The policy is also responsible for deciding whether to de-provision an existing CH at the end of each billing time unit. Additionally, developers can define any number of *monitoring metrics*. Metrics are simple 3-tuples  $\langle name, type, cep-statement \rangle$ . CEP-statements are defined over the stream of monitoring events. An example, which defines a metric `AvgEngineSetupTime` of type `java.lang.Double` as the average duration value of all `EngineSetupEvents` received in a 10 second batch, is given in Listing 2.

```

1 MonitoringMetric metric =
2   new MonitoringMetric();
3   metric.setName("AvgEngineSetupTime");
4   metric.setType(Double.class);
5   metric.setEpl(
6     "select avg(duration)
7     from EngineSetupEvent.win
8     :time_batch(10 sec)"
9   );
10  EventCorrelationEngine.getInstance()
11    .registerMetric(metric);

```

Listing 2. Defining Monitoring Metrics via CEP

*Monitoring metrics* range from very simple and domain-independent (e.g., calculating the average CPU utilization of all CHs) to rather application-specific ones, such as the example given in Listing 2. Whenever the CEP-statement is triggered, the CEP engine writes a new value to an in-memory *monitoring repository*. *Scaling policies*

have access to this repository, and make use of its content in their decisions. In combination with monitoring metrics, scaling policies are a well-suited tool for developers to specify how the application should react to changes in its work load. Hence, sophisticated scaling policies that minimize cloud infrastructure costs or that maximize utilization [Genaud and Gossa 2011] are easy to integrate. As part of the JCLOUDSCALE release, we provide a small number of default policies that users can integrate out of the box, but expect users to write their own policy for non-trivial applications. This has proven necessary as, generally, no generic scaling policy is able to cover the needs of all applications.

Finally, the *cloud manager* component, which can be seen as the heart of the JCLOUDSCALE client-side middleware and the executor of the MAPE loop, enacts the decisions of the policy by invoking the respective functions of the IaaS API and the CH remote interfaces (e.g., provisioning of new CHs, de-provisioning of existing ones, as well as the deployment or migration of COs).

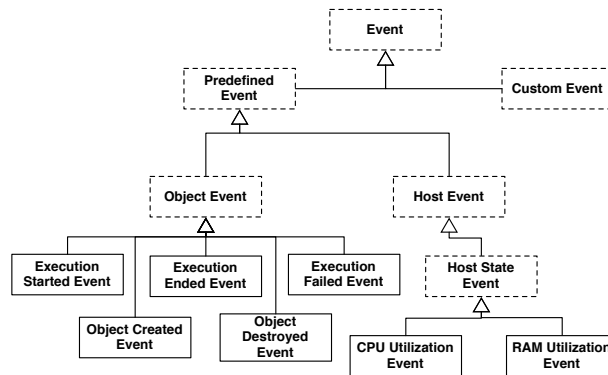


Fig. 4. Monitoring Event Hierarchy

Fig. 4 depicts the type hierarchy of all predefined events in JCLOUDSCALE. Dashed classes denote abstract events, which are not triggered directly, but serve as classifications for groups of related events. All events further contain a varying number of event properties, which form the core information of the event. For instance, for `ExecutionFailedEvent`, the properties contain the CO, the invoked method, and the actual error. Developers and *external event sources* can extend this event hierarchy by inheriting from `CustomEvent`, and writing these custom events into a special event sink (injected by the middleware, see Listing 1). This process is described in more detail in [Leitner et al. 2012].

### 3.2. Deploying to the Cloud

As all code that interacts with the IaaS cloud is injected, the JCLOUDSCALE programming model naturally decouples Java applications from the cloud environment that they are physically deployed to. This allows developers to re-deploy the same application to a different cloud simply by changing the respective parts of the JCLOUDSCALE configuration. JCLOUDSCALE currently contains three separate cloud backends, supporting OpenStack-based private clouds, the Amazon EC2 public cloud, and a special *local environment*. The local environment does not use an actual cloud at all, but simulates CHs by starting new JVMs on the same physical machine as the target appli-

cation. Support for more IaaS clouds, for instance Microsoft Azure’s virtual machine cloud, is an ongoing activity.

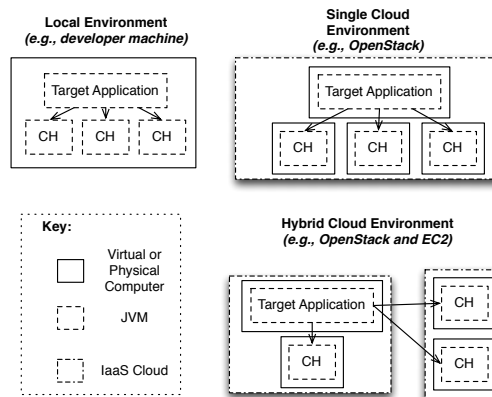


Fig. 5. Supported Deployment Environments

It is also possible to combine different environments, enabling hybrid cloud applications. In this case, the scaling policy is responsible for deciding which CO to execute on which cloud. Fig. 5 illustrates the different types of environments supported by JLOUDSCALE.




Fig. 6. Conceptual Development Process

### 3.3. Development Process

As JLOUDSCALE makes it easy to switch between different cloud environments, the middleware supports a streamlined development process for elastic applications, as sketched in Fig. 6. The developer typically starts by building her application as a local, multi-threaded Java application using common software engineering tools and methodologies. Once the target application logic is implemented and tested, she adds the necessary JLOUDSCALE annotations, as well as scaling policies, monitoring metric definitions, and JLOUDSCALE configuration as required. Now she enables JLOUDSCALE code injection by adding the necessary post-compilation steps to the application build process. Via configuration, the developer specifies a deployment in the local environment first. This allows her to conveniently test and debug the application on her development machine, including tuning and customizing the scaling policy. Finally, once she is satisfied with how the application behaves, she changes the configuration to an actual cloud environment, and deploys and tests the application in a physically distributed fashion.



We argue that this process significantly lessens the pain that developers experience when building applications for IaaS clouds, as it reduces the error-prone and time-consuming testing of applications on an actual cloud. However, of course this process is idealized. Practical usage shows that developers will have to go back to a previous step in the process on occasion. For instance, after testing the scaling behavior in the local environment, the developer may want to slightly adapt the target application to better support physical distribution. Still, we have found that the conceptual separation of target application development and implementation of the scaling behavior is well-received by developers in practice.

#### 4. IMPLEMENTATION

We have implemented JCLLOUDSCALE as a Java-based middleware under an Apache Licence 2.0. The current stable version is available from GitHub<sup>1</sup>. This web site also contains documentation and sample applications to get users started with JCLLOUDSCALE.

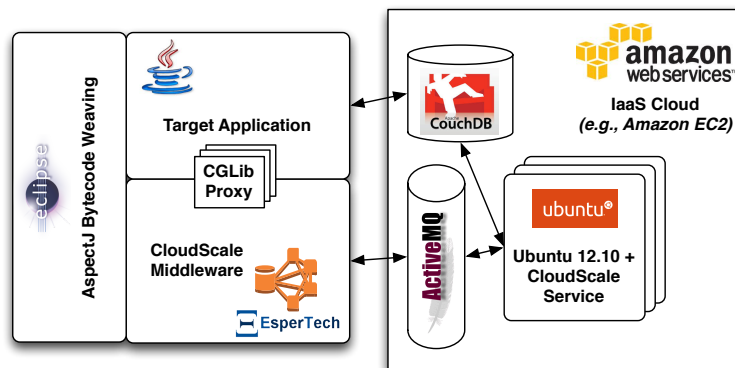


Fig. 7. Implementation Overview

Our implementation integrates a plethora of existing technologies, which is summarized in Fig. 7. JCLLOUDSCALE uses aspect-oriented programming [Kiczales and Hilsdale 2001] (via AspectJ) to inject remoting and cloud management code into target applications. Typically, this is done as a post-compilation step in the build process. Dynamic proxying is implemented by means of the CGLib code generation library. For event processing, the well-established Esper CEP engine is used. The client-side middleware interacts with CHs via a JMS-compatible message queue (currently Apache ActiveMQ). Furthermore, COs and the target application itself can read from and write to a shared data store (for example Apache CouchDB). CHs themselves are simple Ubuntu 12.10 Linux hosts running Java and a small JCLLOUDSCALE operating system service, which receives CO requests and executes them. Currently, we have built CH base images for OpenStack and Amazon EC2, which are linked from the Google Code web site, and which can be used out of the box with the stable version 0.4.0 of JCLLOUDSCALE (the current version at the time of writing). We will also provide images for future stable versions, once they become available.

<sup>1</sup><https://github.com/xLeitix/jcloudscale>

## 5. VALIDATION

As part of our validation of the JCLLOUDSCALE framework, we aim at answering the following three research questions:

- **RQ1:** Does using JCLLOUDSCALE instead of established tooling lead to more efficient development of cloud solutions, e.g., in terms of solution size or development time?
- **RQ2:** How does JCLLOUDSCALE compare with established tooling in terms of ease-of-use, debugging, and other more “soft” quality dimensions?
- **RQ3:** What runtime overhead does JCLLOUDSCALE impose at execution time?

In order to answer **RQ1** and **RQ2**, we conducted a multi-month user study. **RQ3** is addressed via numerical overhead measurements on a simple example application, with and without JCLLOUDSCALE.

### 5.1. User Study

In order to evaluate **RQ1** and **RQ2**, we conducted an user study with 14 participants to assess the developers’ experience with JCLLOUDSCALE as compared to using standard tools.

*5.1.1. Study Setup and Methodology.* We conducted our study with, in total, 14 male master students of computer science at TU Vienna (participants P01 – P14), and based on two different non-trivial implementation tasks. The first task was to develop a parallel computing implementation of a genetic algorithm (**T1**). The second task required the participants to implement a service that executes JUnit test cases on demand (**T2**). For both tasks, an elastic solution was asked for, which was able to react to changes in load dynamically and automatically by scaling up and down in the cloud. Both **T1** and **T2** required roughly a developer week of effort (assuming that the respective participant did not have any particular prior experience with the used technologies).

The study ran in two phases. In Phase (1), we compared using JCLLOUDSCALE on top of OpenStack with programming directly via the OpenStack API, without any specific middleware support. This phase reflected a typical private cloud [Dillon et al. 2010] use case of JCLLOUDSCALE. In Phase (2), we compare JCLLOUDSCALE on top of Amazon EC2 with using Amazon Elastic Beanstalk. This reflects a common public cloud usage of the framework. In both study phases, we asked participating developers to build solutions for both tasks using JCLLOUDSCALE and the respective comparison technology, and compare the developer experience based on quantitative and qualitative factors. We had 9 participating developers in Phase (1), and 5 participants in Phase (2). One participant in Phase (2) only completed one of the two tasks.

Table I. Study Participant Overview

ID	Phase	Java Exp.	Cloud Exp.	JCS/OS	OS	JCS/EC2	Beanstalk
P01	Phase (1)	+	+	T1	T2	–	–
P02	Phase (1)	+	+	T1	T2	–	–
P03	Phase (1)	~	~	T2	T1	–	–
P04	Phase (1)	-	-	T1	T2	–	–
P05	Phase (1)	~	-	T2	T1	–	–
P06	Phase (1)	+	-	T2	T1	–	–
P07	Phase (1)	+	+	T2	T1	–	–
P08	Phase (1)	+	~	T2	T1	–	–
P09	Phase (1)	+	~	T2	T1	–	–
P10	Phase (2)	+	+	–	–	T2	T1
P11	Phase (2)	+	~	–	–	T1	T2
P12	Phase (2)	+	-	–	–	T1	T2
P13	Phase (2)	+	~	–	–	T2	T1
P14	Phase (2)	+	+	–	–	T2	–

Phase (1) of the study lasted two months. We initially presented JCLLOUDSCALE and the comparison technologies to the participants, and randomly assigned which of the tools each participant should be using for T1. Participants then had one month of time to submit a working solution to the task along with a short report, after which they could start working on T2 with the remaining technology. Similar to T1, participants were given one month of time to submit a solution and a short report. Based on the lessons learned from Phase (1), we slightly clarified and improved the task descriptions and gave participants more time (1.5 months per task) for Phase (2). Other than that, Phase (2) was executed identically to Phase (1).

Table I summarizes the relevant background for each participant of the study. To preserve anonymity, we classify the self-reported background of participants related to their Java or cloud experience into three groups: relevant work experience (+), some experience (~), or close to no experience (-). The last four columns indicate whether the participant submitted solutions for JCLLOUDSCALE running on top of OpenStack, OpenStack directly, JCLLOUDSCALE running on top of EC2, or AWS Elastic Beanstalk, as well as which task the participant solved using these (combinations of) technologies.

For the OpenStack-related implementations, we used a private cloud system hosted at TU Vienna. This OpenStack instance consists of 12 dedicated Dell blade servers with 2 Intel Xeon E5620 CPUs (2.4 GHz Quad Cores) each, and 32 GByte RAM, running on OpenStack Folsom (release 2012.2.4). For the study, each participant was allotted a quota of up to 8 very small instances (1 virtual CPU, and 512 MByte of RAM), which they could use to implement and test their solutions. For the AWS-related implementations, participants were assigned an AWS account with sufficient credit to cover their implementation and testing with no particular limitations.

Table II. Solutions Sizes (in Lines of Code)

	Phase (1)				Phase (2)			
	JCS/OS		OS		JCS/EC2		Beanstalk	
	$\tilde{A}$	$\sigma_A$	$\tilde{B}$	$\sigma_B$	$\tilde{C}$	$\sigma_C$	$\tilde{D}$	$\sigma_D$
<b>T1</b>								
Business Logics	<b>200</b>	176	552	215	<b>388</b>	152	825	947
Cloud Management	<b>100</b>	112	180	86	<b>163</b>	24	676	742
Other Code	<b>170</b>	157	286	226	1590	1203	<b>897</b>	1127
Entire Application	<b>400</b>	416	1050	434	<b>2141</b>	1331	2790	2660
<b>T2</b>								
Business Logics	450	292	<b>375</b>	669	800	434	<b>208</b>	280
Cloud Management	<b>100</b>	48	250	364	<b>118</b>	745	223	38
Other Code	325	213	<b>300</b>	297	<b>140</b>	2240	1290	972
Entire Application	<b>1025</b>	461	1500	901	<b>1000</b>	3328	2184	968

5.1.2. *Comparison of Development Efforts (RQ1)*. **RQ1** asked whether JCLLOUDSCALE makes it easier and faster to build elastic IaaS applications. To this end, we asked participants to report on the size of their solutions (in lines of code, without comments and blank lines). The results are summarized in Table II.  $\tilde{A} - \tilde{D}$  represent the median size of solutions, while  $\sigma_A - \sigma_D$  indicate standard deviations. It can be seen that using JCLLOUDSCALE indeed generally reduces the total source code size of applications. Going into the study, we expected JCLLOUDSCALE to mostly reduce the amount of code necessary for interacting with the cloud. However, our results indicate that using JCLLOUDSCALE also often reduced the amount of code of the application business logics, as well as assorted other code (e.g., data structures). When investigating these results, we found that participants considered many of the tasks that JCLLOUDSCALE

takes over as “business logics” when building the elastic application on top of OpenStack or Elastic Beanstalk. To give one example, many participants counted code related to performance monitoring towards “business logics”. Note that, due to the open nature of our study tasks, the standard deviations are all rather large (i.e., solutions using all technologies varied widely in size). It needs to be noted that the large difference in T1 sizes (for JCLLOUDSCALE on top of OpenStack and EC2) between Phase (1) and Phase (2) solutions can be explained by clarifications in the task descriptions. In Phase (1), some formulations in the tasks led to much simpler implementations, while our requirements were formulated much more unambiguously in Phase (2), leading to more complex (and larger) submissions. Hence, we caution the reader to not compare results from Phase (1) with those from Phase (2).

Table III. Time Spent (in Full Hours)

		Phase (1)				Phase (2)			
		JCS/OS		OS		JCS/EC2		Beanstalk	
		$\tilde{A}$	$\sigma_A$	$\tilde{B}$	$\sigma_B$	$\tilde{C}$	$\sigma_C$	$\tilde{D}$	$\sigma_D$
<b>T1</b>									
	Tool Learning	<b>7</b>	2	12	5	28	18	<b>16</b>	1
	Coding	<b>4</b>	10	30	17	<b>42</b>	25	54	23
	Bug Fixing	<b>7</b>	7	18	12	<b>14</b>	8	20	14
	Other Activities	<b>13</b>	9	14	14	<b>5</b>	0	6	6
	Entire Application	<b>31</b>	25	76	33	127	25	<b>121</b>	36
<b>T2</b>									
	Tool Learning	8	6	<b>2</b>	1	<b>15</b>	10	23	18
	Coding	30	11	<b>25</b>	17	36	5	<b>30</b>	14
	Bug Fixing	10	11	10	7	16	16	<b>5</b>	0
	Other Activities	<b>7</b>	4	11	10	<b>5</b>	0	9	9
	Entire Application	62	13	<b>46</b>	17	125	40	<b>102</b>	13

However, looking at lines of code alone is not sufficient to validate our hypothesis that JCLLOUDSCALE improves developer productivity, as it would be possible that the JCLLOUDSCALE solutions, while being more compact, are also more complicated (and, hence, take longer to implement). That is why we also asked participants to report on the time they spent working on their solutions. The results are compiled in Table III. We have classified work hours into a number of different activities: initially learning the technology, coding, testing and bug fixing, and other activities (e.g., building OpenStack cloud images). Our results indicate that the initial learning curve for JCLLOUDSCALE is lower than for working with OpenStack directly. However, in comparison with Elastic Beanstalk, some participants reported equal or even more complexity of JCLLOUDSCALE, mainly because of the limited tutorial and help information about JCLLOUDSCALE available in Internet. For coding JCLLOUDSCALE appeared to be much faster tool for participants who had at least some prior experience with cloud computing.

Due to the high standard deviations, looking at this quantitative data alone remains inconclusive. Hence, we also analyzed qualitative feedback by the participants in their reports. Multiple developers have reported that they felt more productive when using JCLLOUDSCALE. For instance, P01 has stated that “*the coolest thing about JCLLOUDSCALE is the reduction of development effort necessary, to host applications in the cloud (...) [there] are a lot of thing you do not have to care about in detail.*” P03 also concluded that using JCLLOUDSCALE “*went a lot smoother than [using OpenStack directly]*”. P07 also seemed to share this sentiment and stated that “*[After resolving initial problems] the rest of the project was without big problems and I was able to be very productive in coding the solution.*” In comparison to Elastic Beanstalk, participants indicate that core idea behind JCLLOUDSCALE is easier to grasp for starting cloud developers than

the one behind modern PaaS systems. For example, P13 indicated that “*the [ JCLLOUDSCALE ] API is easier to understand and more intuitive to use. Also it fits more into a Java-like programming model, instead of the weird request based approach of the amazon API*”. However, some participants noted that the fact that Elastic Beanstalk is based on common technology also appeals to them. For instance, P10 specified that “*[In case of Elastic Beanstalk,] Well-known technology is the basis for everything (Tomcat/Servlet)*”. Hence, the participant argued that this allows developers who are already familiar with these platforms to be productive sooner.

Summarizing our study results regarding **RQ1**, we conclude that JCLLOUDSCALE indeed seems to allow for higher developer productivity than working directly on top of an IaaS API, such as the one of OpenStack. In comparison to AWS Elastic Beanstalk, our results do not clearly indicate more or less productivity with JCLLOUDSCALE.

5.1.3. *Comparisong of Developer-Perceived Quality (RQ2)*. In order to answer **RQ2**, we were interested in the participant’s subjective evaluation of the used technologies. Hence, we asked them to rate the technologies along a number of dimensions from 1 (very good) to 5 (insufficient). We report on the dimensions “simplicity” (how easy is it to use the tool?), “debugging” (how easy is testing and debugging the application?), “development process” (does the technology imply an awkward development process?), and “stability” (how often do unexpected errors occur?). A summary of our results is shown in Table IV.

Table IV. Subjective Ratings (lower is better)

		Phase (1)				Phase (2)			
		JCS on OpenStack		OpenStack		JCS on EC2		Elastic Beanstalk	
		$\bar{A}$	$\sigma_A$	$\bar{B}$	$\sigma_B$	$\bar{C}$	$\sigma_C$	$\bar{D}$	$\sigma_D$
<b>T1</b>	Simplicity	3	0.6	3	1.2	2	0	2	1.4
	Debugging	3	1.5	3	1	4	0	<b>3.5</b>	0.7
	Development Process	4	1.7	<b>3.5</b>	0.5	<b>2</b>	1.4	3	0
	Stability	2	1.4	2	0.8	2	1.4	<b>1.5</b>	0.7
	Overall	3	0.6	3	0.8	2	0	2	1.4
	<b>T2</b>								
	Simplicity	<b>2</b>	0.4	3	1.4	<b>2</b>	1.5	3	1.4
	Debugging	<b>2</b>	0.7	4	1.4	4	0	4	0
	Development Process	<b>2</b>	0.6	3	1.4	<b>2</b>	0	2.5	0.7
	Stability	2	1.5	<b>1</b>	0	3	0.5	<b>2.5</b>	0.7
	Overall	<b>2</b>	0.4	3	0	3	0.5	3	1.4

Overall, participants rated all used technologies similarly. However, JCLLOUDSCALE was rated worse than the comparison technologies mainly in terms of “stability”. This is not a surprise, as JCLLOUDSCALE still is a research prototype in a relatively early development stage. Participants indeed mentioned multiple stability-related issues in their reports (e.g., P10 mentions that “*When deploying many cloud objects to one host there were behaviors which were hard to reason about*”). Further, some technical implementation decisions in JCLLOUDSCALE were not appreciated by our study participants. To give an example, P11 noted that “*It is confusing in the configuration that the field AMI-ID actually expects the AMI-Name, not the ID*”. In contrast, JCLLOUDSCALE has generally been rated slightly better in terms of simplicity and ease-of use, especially for T2. For example, participant P09 claimed that “*JCLLOUDSCALE is the clear winner in ease of use. If you quickly want to just throw some Objects in the cloud, it’s the clear choice.*”. Similarly, P12 reported “*[JCLLOUDSCALE is] programmer friendly. All procedure is more low level and as a programmer there are more things to tune and adjust.*”. In terms of debugging features, all used technologies were not rated overly

well. JCLLOUDSCALE was generally perceived slightly better (due to its local development environment), but realistically all compared systems are currently deemed too hard to debug if something goes wrong. Finally, in terms of the associated development process, JCLLOUDSCALE is generally valued highly, with the exception of T1 and JCLLOUDSCALE on Openstack. We assume that this is a statistical artifact, as the development process of JCLLOUDSCALE is judged well in all other cases. Concretely P01 stated that with JCLLOUDSCALE, *“You are able to get application into the cloud really fast. You are not forced to take care about a lot of cloud-specific issues”*.

Independently of the subjective ratings, multiple participants stated that they valued the flexibility that the JCLLOUDSCALE concept brought over Elastic Beanstalk. Particularly, P11 noted that *“[JCLLOUDSCALE provides] more flexibility. The developer can decide when to deploy hosts, on which host an object gets deployed, when to destroy a host, etc”*. Additionally, participants favored the monitoring event engine of JCLLOUDSCALE for performance tracking over the respective features of the PaaS system. For example, P12 specified as an JCLLOUDSCALE advantage that *“programmatic usage of different events with a powerful event correlation framework [is] in combination with listeners extremely powerful.”*

Concluding our discussion regarding **RQ2**, we note that JCLLOUDSCALE indeed has some way to go before it is ready for industrial usage. The general concepts of the tool are valued by developers, but currently, technical issues and lack of documentation and technical support make it hard for developers to fully appreciate the power of the JCLLOUDSCALE model. One aspect that needs more work is how developers define the scaling behavior of their application. Both tasks in our study required the participants to define non-trivial scaling policies, e.g., in order to optimally schedule genetic algorithm executions to cloud resources, which most participants felt unable to do with the current API provided by JCLLOUDSCALE. Overall, in comparison to working directly on OpenStack, many participants preferred JCLLOUDSCALE, but compared to a mature PaaS platform, AWS Elastic Beanstalk still seems slightly preferable to many. However, it should be noted that JCLLOUDSCALE still opens up use cases for which using Beanstalk is not an option, for instance for deploying applications in a private or hybrid cloud [Leitner et al. 2013].

## 5.2. Runtime Overhead Measurements (RQ3)

Finally, we investigated whether the improved convenience of JCLLOUDSCALE is paid for with significantly reduced application performance. Therefore, the main goal of these experiments was to compare the performance of the same application built on top of JCLLOUDSCALE and using IaaS platform (OpenStack or EC2) directly.

*5.2.1. Experiment Setup.* To achieve this, we built a simple sample application (inspired by T2 from the user study) on top of Amazon EC2 and our private OpenStack cloud. The application JSTAAS (“JavaScript Testing-as-a-Service”) provides testing of JavaScript applications as a cloud service. Clients register with the service, which triggers JSTAAS to periodically launch this client’s registered test suites. Results of test runs are stored in a database, which can be queried by the client. Tests vary widely in the load that they generate on the servers, and clients are billed according to this load.

Secondly, we also implemented the same application using JCLLOUDSCALE. As the main goal was to calculate the overhead introduced by the JCLLOUDSCALE, we designed both implementations to have the same behavior and reuse as much business logics code as possible. In addition, to simplify our setup, focus on execution performance evaluation and to avoid major platform-dependent side effects, we limited our-




Fig. 8. Execution on OpenStack Platform

selves to a scenario, where the number of available cloud hosts is static. The source code of both applications is available online<sup>2</sup>.

All four solutions (directly on OpenStack, directly on EC2, and using JCLLOUDSCALE on both, OpenStack and EC2) follow a simple master-slave pattern: a single node (the master) receives tests through a SOAP-based Web Service and schedules them over the set of available worker nodes. All solutions were tested with a test setup that consisted of 40 identical parallelizable long-running test suites scheduled evenly over the set of available cloud machines. Each test suite consisted of a set of dummy JavaScript tests calculating Fibonacci numbers. During the evaluation, we measured the total time of an entire invocation of the service (i.e., how long a test request takes end-to-end, including scheduling, data transmission, result collection, etc.). A single experiment run consisted of 10 identical invocations of the testing web service, each time with a different number of CHs (ranging from 2 to 20 CHs). To eliminate the influence of other parallel cloud activity, experiment runs for both versions of evaluation application were running interchangeably for each number of CHs. To discard accidental results, such experiment runs were repeated 10 times for each version of evaluation application.

**5.2.2. Experiment Results.** Figure 8 and Figure 9 show the median total execution time for different numbers of hosts. In general, both applications show similar behavior in each environment, meaning that both approaches are feasible and have similar parallelizing capabilities with minor overhead difference. In both environments, there is an overhead of JCLLOUDSCALE that is proportional to the amount of used CHs and approximately equal to 2 – 3 seconds per introduced host for multiple minutes evaluation application. This overhead may be significant for performance-critical production applications, but we believe that it is a reasonable price to pay in the current development stage of the JCLLOUDSCALE middleware. However, detailed investigation (and, subsequently, reduction) of the overhead introduced by JCLLOUDSCALE is planned for future releases of the system.

Another important issue that is visible from Figure 8 and Figure 9 is the cloud performance stability and predictability. With an increasing number of hosts, the total execution time is expected to monotonously decrease, up to a limit when the overhead of parallelization is larger than the gain of having more processors available. This indeed happens in case of Amazon EC2. However, starting with 10 used hosts in OpenStack, the overall application execution time remains almost constant or even increases. In our case, this is mainly caused by the limited size of our private cloud.

<sup>2</sup><http://www.infosys.tuwien.ac.at/staff/phdschool/rstzab/papers/TOIT14/>




Fig. 9. Execution on EC2 Platform

In our system, starting with 10 hosts, physical machines start to get overutilized, and virtual machines start to compete for resources (e.g., CPU or disk IO).

### 5.3. Threats to Validity

The major threat to (internal) validity, which the results relating to **RQ1** and **RQ2** face, are that the small sample size of 14 study participants, along with relatively open problem statements, does not allow us to establish statistical significance. However, due to the reports we received from participants, as well as due to comparing the solutions themselves, we are convinced that our results that **JCLOUDSCALE** lets developers build cloud applications more efficiently was not a coincidence. Further, our participants were aware that **JCLOUDSCALE** is our own system. Hence, there is a chance that our participants gave their reports a more positive spin, so as to do us a favor. However, given that all reports contained both negative aspects of all evaluated frameworks, we are confident that most participants reported truthfully. In terms of external validity, it is possible that the two example projects we chose for our study are not representative of real-world applications. However, we argue that this is unlikely, as the projects have specifically been chosen based on real-life examples that the authors of this paper are aware of or had to build themselves in the past. Another thread to external validity is that the participants of our study are all students at TU Vienna. While most of them have some practical real-life experience in application development, none can be considered senior developers.

In terms of **RQ3**, the major threat to external validity is that the application we used to measure overhead on is necessarily simplified, and not guaranteed to be representative of real **JCLOUDSCALE** applications. Real applications, such as the ones built in our user study, are hard to replicate in exactly the same way on different systems, hence comparative measurements amongst such systems are always unfair. To minimize this risk, we have taken care to preserve what we consider core features of cloud applications even in the simplified measurement application.

## 6. RELATED WORK

We now put the **JCLOUDSCALE** framework into context of the larger distributed and cloud computing ecosystem. As the scope of **JCLOUDSCALE** is rather wide, there are a plethora of existing tools, frameworks and middleware that are related to parts of the functionality of our system. Based on the descriptions in Section 2 and Section 3, we consider the main dimensions to compare frameworks across are (1) to what extent they transparently handle remoting and elasticity, (2) how easy it is to locally test and debug applications, (3) whether the system restricts what kinds of applications can be built (e.g., only Tomcat-based web applications), (4) whether the system handles cloud



	Transparent Remoting	Transparent Elasticity	Local Testing	Unrestricted Architecture	Transparent VM Management	Cloud Provider Independence
<b>Remoting Frameworks</b>						
Java RMI	yes	no	yes	yes	no	no
Enterprise Java Beans (EJB)	yes	partial	yes	yes	no	no
Elastic Remote Methods	yes	yes	no	yes	yes	yes
Aneka [Vecchiola et al. 2008]	partial	no	no	yes	yes	yes
[Calheiros et al. 2012]						
<b>PaaS Systems</b>						
Appengine	yes	yes	partial	no	yes	no
Amazon Elastic Beanstalk	yes	yes	no	no	yes	no
Heroku	yes	yes	partial	no	yes	no
AppScale [Chohan et al. 2010]	yes	yes	no	no	yes	yes
ConPaaS [Krintz 2013]	yes	yes	no	partial	yes	yes
[Pierre et al. 2011]						
[Pierre and Stratan 2012]						
BOOM [Alvaro et al. 2010]	yes	yes	no	no	no	yes
Esc [Satzger et al. 2011]	yes	yes	no	no	no	yes
Granules [Pallickara et al. 2009]	yes	yes	yes	no	no	yes
<b>Cloud Deployment &amp; Test Frameworks</b>						
JClouds	no	no	no	yes	no	yes
Docker	no	no	yes	yes	no	yes
Cafe [Mietzner et al. 2009]	no	no	yes	no	no	yes
MADCAT [Inzinger et al. 2014]	no	no	yes	no	no	yes
OpenTOSCA [Binz et al. 2012]	no	partial	no	yes	yes	yes
[Binz et al. 2013]						
JCLOUDSCALE	yes	partial	yes	yes	yes	yes

virtual machines for the user, and (5) whether the system is bound to one specific cloud provider. Systems that are cloud provider independent typically also can be used in a private cloud context. We provide a high-level comparison of various systems along these axes in Table 6. We have evaluated each system along each axis, and assigned “yes” (fully supported), “no” (no real support), or “partial” (some support). All evaluations are to the best of the knowledge of the authors of this paper, and based on tool documentations or publications.

Firstly, JCLLOUDSCALE can be compared to traditional distributed object middleware [Wolfgang Emmerich 2000], such as Java RMI or EJB. These systems provide transparent remoting features, comparable to JCLLOUDSCALE, but clearly do not provide any support for cloud specifics, such as VM management. It can be argued that EJB provides some amount of transparent elasticity, as EJB application containers can be clustered. However, it is not easy to scale such clusters up and down. A recent research work [Jayaram 2013] has introduced the idea of Elastic Remote Methods, which extends Java RMI with cloud-specific features. This work is comparable in goals to our contribution. However, the technical approach is quite different. Aneka [Vecchiola et al. 2008; Calheiros et al. 2012], a well-known .NET based cloud framework, is a special case of a cloud computing middleware that also exhibits a number of characteristics of a PaaS system. We argue that Aneka’s abstraction of remoting is not perfect, as developers are still rather intimately aware of the distributed processing that is going on. To the best of our knowledge, Aneka does not automatically scale systems, and provides no local testing environment.

Secondly, as already argued in Section 5, many of JCLLOUDSCALE’s features are comparable to common PaaS systems (Google Appengine, Amazon Elastic Beanstalk, or Heroku, to name just a few). All of these platforms provide transparent remoting and elasticity, and take over virtual machine management from the user. However, all of these systems also require a relatively specific application architecture (usually a three-tiered web application), and usually tie the user tightly to one specific cloud provider. Support for local testing is usually limited, although most providers nowadays have at least limited tooling or emulators available for download.

In addition to these commercial PaaS systems, there are also multiple platforms coming out of a research setting. For instance, AppScale [Chohan et al. 2010; Krintz 2013] is an open-source implementation of the Google Appengine model. AppScale can also be deployed on any IaaS system, making it much more vendor-independent than other PaaS platforms. This is similar to the ConPaaS open source platform [Pierre et al. 2011; Pierre and Stratan 2012], which originates from a European research project of the same name. ConPaaS follows a more service-oriented style, treating applications as collections of loosely-coupled services. This makes ConPaaS suited for a wider variety of applications, however, it has to be said that ConPaaS still imposes significant restrictions on the application that is to be hosted.

In scientific literature, there are also a number of PaaS systems which are more geared towards data processing, e.g., BOOM [Alvaro et al. 2010], Esc [Satzger et al. 2011], or Granules [Pallickara et al. 2009]. These systems are hard to compare with our work, as they generally operate in an entirely different fashion as compared to JCLLOUDSCALE or the commercial PaaS operators. However, they typically only support a very restricted type of (data-driven) application model, and often do not actually interact with the cloud by themselves. This makes them necessary cloud provider independent, but also means that developers need to implement the actual elasticity-related features themselves.

Thirdly, we need to compare JCLLOUDSCALE to a number of cloud computing related frameworks, which cover a part of the functionality provided by our middleware. JClouds is a Java library that abstracts from the heterogeneous APIs of different IaaS

providers, and allows to decouple Java applications from the IaaS system that they operate in. JCLLOUDSCALE internally uses JClouds to interact with providers. However, by itself, JClouds does not provide any actual elasticity. Docker is a container framework geared towards bringing testability to cloud computing. Essentially, Docker has similar goals to the local test environment of JCLLOUDSCALE.

JCLLOUDSCALE also has some relation to the various cloud deployment models and systems that have recently been proposed in literature, e.g., Café [Mietzner et al. 2009], MADCAT [Inzinger et al. 2014], or OpenTOSCA [Binz et al. 2012; Binz et al. 2013], which is an open source implementation of an upcoming OASIS standard. These systems do not typically cover elasticity by themselves (although TOSCA has partial support for auto-scaling groups), but they are usually independent of any concrete cloud provider.

By design, JCLLOUDSCALE supports most of the characteristics we discuss here. However, especially in comparison to PaaS systems, developers of JCLLOUDSCALE applications are not entirely shielded from issues of scalability. Further, as the user study discussed in Section 5 has shown, the system still needs to improve how scaling policies are written, so as to make building elastic systems easier for developers.

## 7. CONCLUSIONS

JCLLOUDSCALE is a Java-based middleware that eases the development of elastic cloud applications on top of an IaaS cloud. JCLLOUDSCALE follows a declarative approach based on Java annotations, which removes the need to actually adapt the business logics of the target application to use the middleware. Hence, JCLLOUDSCALE support can easily be turned on and off for an application, leading to a flexible development process that clearly separates the implementation of target application business logics from implementing and tuning the scaling behavior.

We have introduced the core concepts behind JCLLOUDSCALE, and presented an evaluation of the middleware based on an user study as well as using a simple case study application. Our results indicate that JCLLOUDSCALE is well received among initial developers. Our results support our claim that the general JCLLOUDSCALE model has advantages to both, working directly on top of an IaaS API or on an industrial PaaS systems. However, further study is required to strengthen these claims, as the limited scale of our initial study was not sufficient to clear all doubts about the viability of the system. Further, there are also technical and conceptual issues that require further investigation. Most importantly, we have learned that implementing actually elastic applications is still cumbersome for developers, as getting the scaling policy right is not as easy as we had hoped.

## REFERENCES

- Linton Abraham, Michael A. Murphy, Michael Fenn, and Sebastien Goasguen. 2010. Self-Provisioned Hybrid Clouds. In *Proceedings of the 7th International Conference on Autonomic Computing (ICAC'10)*. ACM, New York, NY, USA, 161–168. DOI: <http://dx.doi.org/10.1145/1809049.1809075>
- Peter Alvaro, Tyson Condie, Neil Conway, Khaled Elmeleegy, Joseph M. Hellerstein, and Russell Sears. 2010. Boom Analytics: Exploring Data-Centric, Declarative Programming for the Cloud. In *Proceedings of the 5th European Conference on Computer Systems (EuroSys'10)*. ACM, 223–236.
- Michael Armbrust, Armando Fox, Rean Griffith, Anthony D Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, Ariel Rabkin, Ion Stoica, and Matei Zaharia. 2010. A View of Cloud Computing. *Communications of the ACM* 53, 4 (2010), 50–58.
- Cor-Paul Bezemer, Andy Zaidman, Bart Platzbeecker, Toine Hurkmans, and Aad 't Hart. 2010. Enabling Multi-Tenancy: An Industrial Experience Report. In *Proceedings of the 2010 IEEE International Conference on Software Maintenance (ICSM '10)*. IEEE Computer Society, Washington, DC, USA, 1–8. DOI: <http://dx.doi.org/10.1109/ICSM.2010.5609735>

- Tobias Binz, Uwe Breitenbücher, Florian Haupt, Oliver Kopp, Frank Leymann, Alexander Nowak, and Sebastian Wagner. 2013. OpenTOSCA - A Runtime for TOSCA-Based Cloud Applications. In *Proceedings of the 11th International Conference on Service-Oriented Computing (ICSOC)*. 692–695.
- Tobias Binz, Gerd Breiter, Frank Leyman, and Thomas Spatzier. 2012. Portable Cloud Services Using TOSCA. *IEEE Internet Computing* 16, 3 (May 2012), 80–85. DOI : <http://dx.doi.org/10.1109/MIC.2012.43>
- Rajkumar Buyya, Chee Shin Yeo, Srikumar Venugopal, James Broberg, and Ivona Brandic. 2009. Cloud Computing and Emerging IT Platforms: Vision, Hype, and Reality for Delivering Computing as the 5th Utility. *Future Generation Computing Systems* 25 (2009), 599–616. Issue 6.
- Rodrigo N. Calheiros, Christian Vecchiola, Dileban Karunamoorthy, and Rajkumar Buyya. 2012. The Aneka platform and QoS-driven resource provisioning for elastic applications on hybrid Clouds. *Future Generation Computer Systems* 28, 6 (June 2012), 861–870. DOI : <http://dx.doi.org/10.1016/j.future.2011.07.005>
- Navraj Chohan, Chris Bunch, Sydney Pang, Chandra Krintz, Nagy Mostafa, Sunil Soman, and Rich Wolski. 2010. AppScale: Scalable and Open AppEngine Application Development and Deployment. In *Cloud Computing*, Dimiter Avresky, Michel Diaz, Arndt Bode, Bruno Ciciani, and Eliezer Dekel (Eds.). Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, Vol. 34. Springer Berlin Heidelberg, 57–70. DOI : [http://dx.doi.org/10.1007/978-3-642-12636-9\\_4](http://dx.doi.org/10.1007/978-3-642-12636-9_4)
- Tharam Dillon, Chen Wu, and Elizabeth Chang. 2010. Cloud Computing: Issues and Challenges. In *Proceedings of the 2010 24th IEEE International Conference on Advanced Information Networking and Applications (AINA'10)*. IEEE Computer Society, Washington, DC, USA, 27–33. DOI : <http://dx.doi.org/10.1109/AINA.2010.187>
- Stephane Genaud and Julien Gossa. 2011. Cost-Wait Trade-Offs in Client-Side Resource Provisioning with Elastic Clouds. In *Proceedings of the 2011 IEEE 4th International Conference on Cloud Computing (CLOUD '11)*. IEEE Computer Society, Washington, DC, USA, 1–8. DOI : <http://dx.doi.org/10.1109/CLOUD.2011.23>
- Robert L. Grossman. 2009. The Case for Cloud Computing. *IT Professional* 11, 2 (2009), 23–27. DOI : <http://dx.doi.org/10.1109/MITP.2009.40>
- Christian Inzinger, Stefan Nastic, Sanjin Sehic, Michael Vögler, Fei Li, and Schahram Dustdar. 2014. MADCAT: A Methodology for Architecture and Deployment of Cloud Application Topologies. In *Proceedings of the 8th International Symposium on Service Oriented System Engineering (SOSE)*. 13–22. DOI : <http://dx.doi.org/10.1109/SOSE.2014.9>
- K.R. Jayaram. 2013. Elastic Remote Methods. In *Proceedings of Middleware 2013 (Lecture Notes in Computer Science)*, David Eyers and Karsten Schwan (Eds.), Vol. 8275. Springer Berlin Heidelberg, 143–162. DOI : [http://dx.doi.org/10.1007/978-3-642-45065-5\\_8](http://dx.doi.org/10.1007/978-3-642-45065-5_8)
- Jeffrey O. Kephart and David M. Chess. 2003. The Vision of Autonomic Computing. *Computer* 36, 1 (Jan. 2003), 41–50. DOI : <http://dx.doi.org/10.1109/MC.2003.1160055>
- Gregor Kiczales and Erik Hilsdale. 2001. Aspect-Oriented Programming. *SIGSOFT Software Engineering Notes* 26, 5 (Sept. 2001), 313–. DOI : <http://dx.doi.org/10.1145/503271.503260>
- Chandra Krintz. 2013. The AppScale Cloud Platform: Enabling Portable, Scalable Web Application Deployment. *IEEE Internet Computing* 17, 2 (March 2013), 72–75. DOI : <http://dx.doi.org/10.1109/MIC.2013.38>
- George Lawton. 2008. Developing Software Online With Platform-as-a-Service Technology. *Computer* 41, 6 (June 2008), 13–15. DOI : <http://dx.doi.org/10.1109/MC.2008.185>
- Philipp Leitner, Christian Inzinger, Waldemar Hummer, Benjamin Satzger, and Schahram Dustdar. 2012. Application-Level Performance Monitoring of Cloud Services Based on the Complex Event Processing Paradigm. In *Proceedings of the 2012 5th IEEE International Conference on Service-Oriented Computing and Applications (SOCA)*. IEEE Computer Society, Washington, DC, USA, 1–8. DOI : <http://dx.doi.org/10.1109/SOCA.2012.6449437>
- Philipp Leitner, Zabolotnyi Rostyslav, Alessio Gambi, and Schahram Dustdar. 2013. A Framework and Middleware for Application-Level Cloud Bursting on Top of Infrastructure-as-a-Service Clouds. In *Proceedings of the 2013 IEEE/ACM 6th International Conference on Utility and Cloud Computing (UCC '13)*. IEEE Computer Society, Washington, DC, USA, 163–170. DOI : <http://dx.doi.org/10.1109/UCC.2013.39>
- Philipp Leitner, Benjamin Satzger, Waldemar Hummer, Christian Inzinger, and Schahram Dustdar. 2012. CloudScale: a Novel Middleware for Building Transparently Scaling Cloud Applications. In *Proceedings of the 27th Annual ACM Symposium on Applied Computing (SAC'12)*. ACM, New York, NY, USA, 434–440. DOI : <http://dx.doi.org/10.1145/2245276.2245360>
- David Luckham. 2002. *The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems*. Addison-Wesley Professional.
- Ralph Mietzner, Tobias Unger, and Frank Leymann. 2009. Cafe: A Generic Configurable Customizable Composite Cloud Application Framework. In *On the Move to Meaningful Internet Systems (OTM 2009)*,

- Robert Meersman, Tharam Dillon, and Pilar Herrero (Eds.). Vol. 5870. Springer Berlin / Heidelberg, 357–364. [http://dx.doi.org/10.1007/978-3-642-05148-7\\_24](http://dx.doi.org/10.1007/978-3-642-05148-7_24)
- Shrideep Pallickara, Jaliya Ekanayake, and Geoffrey Fox. 2009. Granules: A Lightweight Streaming Runtime for Cloud Computing With Support for Map-Reduce. In *Proceedings of the IEEE International Conference on Cluster Computing and Workshops (CLUSTER'09)*. IEEE, 1–10.
- Guillaume Pierre, Ismail El Helw, Corina Stratan, Ana Oprescu, Thilo Kielmann, Thorsten Schütt, Jan Stender, Matej Artač, and Aleš Černivec. 2011. ConPaaS: An Integrated Runtime Environment for Elastic Cloud Applications. In *Proceedings of the Workshop, Posters and Demos Track (Middleware'11)*. ACM, New York, NY, USA, Article 5, 2 pages. DOI: <http://dx.doi.org/10.1145/2088960.2088965>
- Guillaume Pierre and Corina Stratan. 2012. ConPaaS: A Platform for Hosting Elastic Cloud Applications. *IEEE Internet Computing* 16, 5 (2012), 88–92. DOI: <http://dx.doi.org/10.1109/MIC.2012.105>
- Benjamin Satzger, Waldemar Hummer, Philipp Leitner, and Schahram Dustdar. 2011. Esc: Towards an Elastic Stream Computing Platform for the Cloud. In *Proceedings of the IEEE Fifth International Conference on Cloud Computing (CLOUD'11), Application and Experience Track*. IEEE Computer Society, Los Alamitos, CA, USA, 348–355. DOI: <http://dx.doi.org/10.1109/CLOUD.2011.27>
- Borja Sotomayor, Rubén S. Montero, Ignacio M. Llorente, and Ian Foster. 2009. Virtual Infrastructure Management in Private and Hybrid Clouds. *IEEE Internet Computing* 13, 5 (Sept. 2009), 14–22. DOI: <http://dx.doi.org/10.1109/MIC.2009.119>
- Christian Vecchiola, Xingchen Chu, and Rajkumar Buyya. 2008. Aneka: a Software Platform for .NET based Cloud Computing. In *Proceedings of the High Performance Computing Workshop*. 267–295.
- Wolfgang Emmerich. 2000. *Engineering Distributed Objects*. John Wiley & Sons.
- Rostyslav Zabolotnyi, Philipp Leitner, and Schahram Dustdar. 2013. Dynamic Program Code Distribution in Infrastructure-as-a-Service Clouds. In *Proceedings of the 5th International Workshop on Principles of Engineering Service-Oriented Systems (PESOS 2013), co-located with ICSE 2013*.