

A Monitoring Infrastructure for the Quality Assessment of Cloud Services

Priscila Cedillo, Javier Gonzalez-Huerta, Silvia Abrahao and Emilio Insfran

Abstract Service Level Agreements (SLAs) specify the strict terms under which cloud services must be provided. The assessment of the quality of services being provided is critical for both clients and service providers. In this context, stakeholders must be capable of monitoring services delivered as Software as a Service (SaaS) at runtime and of reporting any eventual non-compliance with SLAs in a comprehensive and flexible manner. In this paper, we present the definition of an SLA compliance monitoring infrastructure, which is based on the use of models@run.time, its main components and artifacts, and the interactions among them. We place emphasis on the configuration of the artifacts that will enable the monitoring, and we present a prototype that can be used to perform this monitoring. The feasibility of our proposal is illustrated by means of a case study, which shows the use of the components and artifacts in the infrastructure and the configuration of a specific plan with which to monitor the services deployed on the Microsoft Azure© platform.

A prior version of this paper has been published in the ISD2015 Proceedings (<http://aisel.aisnet.org/isd2014/proceedings2015/>).

P. Cedillo · S. Abrahao · E. Insfran
Department of Computer Systems and Computation,
Universitat Politècnica de València, Valencia, Spain
e-mail: sabrahao@dsic.upv.es

E. Insfran
e-mail: einsfran@dsic.upv.es

P. Cedillo (✉)
Department of Computer Science, Faculty of Engineering,
University of Cuenca, Cuenca, Ecuador
e-mail: priscila.cedillo@ucuenca.edu.ec; icedillo@dsic.upv.es

J. Gonzalez-Huerta
Département d'Informatique, Université du Québec à Montréal,
Montreal, Canada
e-mail: gonzalez_huerta.javier@uqam.ca

Keywords Model driven engineering · Models@run.time · Quality assessment · Cloud services · Service level agreements · Software as a service

1 Introduction

Software as a Service (SaaS) has emerged as a software deployment model in recent years that makes software available entirely through the use of a web browser, while hiding the details regarding where the software is hosted or its underlying architecture [1]. SaaS is increasingly being used by web-based applications owing the benefits it provides for both users and service providers [2]. The terms under which a SaaS application is provided must be expressed by using Service Level Agreements (SLAs). Each service is typically accompanied by an SLA that defines the minimal guarantees that the cloud provider offers to its customers [3] (e.g. ensuring the availability of a service at least 99.5 % of the time). Service providers are becoming interested in monitoring cloud services in order to assess compliance with the SLA, thus avoiding possible penalizations and improving service quality [4]. On the customer side, service monitoring provides information and Key Performance Indicators (KPIs) that are useful in the decision-making process [5].

Traditional monitoring technologies are restricted to static and homogeneous environments, and cannot therefore be appropriately applied to cloud environments [6]. Cloud computing has led to the emergence of new issues, challenges, and needs as regards measuring quality (e.g. elasticity, scalability, adaptability, timeliness) [5]. Moreover, when compared with other distributed systems such as Grid Computing, the monitoring of a cloud is more complex because of the differences in both the trust model and the view of resources/services presented to the user [7], in addition to the presence of multiple layers and service paradigms [5]. Unfortunately, existing cloud and general purpose monitoring solutions have several limitations, as reported by Muller et al. [8]: the SLAs they support are not sufficiently expressive to model real-world scenarios. They couple the monitoring configuration with a given SLA specification, the explanations of the violations are difficult to understand and even potentially inaccurate, and some proposals either do not provide an architecture or the cohesion of their elements is low. Furthermore, it is important to have flexible quality monitoring infrastructures that will allow service providers to modify the non-functional requirements (NFRs) to be monitored, based on SLAs variations.

We believe that Model Driven Engineering (MDE) may be a solution as regards providing the flexibility required to monitor infrastructures. However, establishing all the NFRs to be monitored when designing the monitoring infrastructure is not always possible (e.g., owing to SLA renegotiations, the addition of new NFRs to be monitored, changes in the cloud platform). In this context, Baresi and Ghezzi [9] advocate that future software engineering research should be focused on providing software with intelligent support at runtime, thus breaking across the current rigid boundary between development-time and runtime. It is therefore necessary to define approaches that will allow cloud services to be monitored and will also permit the addition of new requirements or the modification of existing ones at runtime

without interrupting the service execution. This challenge can be confronted by using `models@run.time` [9]. A `model@run.time` is employed at runtime in a system and its encoding enables its processing at runtime [10]. Besides, a `model@run.time` is causally-connected to the running system, meaning that a change in this model triggers a corresponding change in the running system and/or vice versa [10].

In a previous paper, we presented the definition of a monitoring process for cloud services by using `models@run.time` [11], in which we established the tasks involved in the monitoring process. In this paper, we extend that work by presenting the monitoring infrastructure that using `models@run.time` is able to: (i) retrieve data from the cloud services during their execution; (ii) calculate derived metrics based on these data; and (iii) report any eventual SLA violations. The contribution of this paper is therefore the definition of a monitoring infrastructure, its main components and the artifacts used by the Monitoring Configurator (i.e., quality meta-models with which to generate the Requirements Quality Model, the SaaS Quality Model and the Runtime Quality Model), along with the interactions among them. The feasibility of our proposal is illustrated through a case study, which shows the use of the components and artifacts involved in the infrastructure and the configuration of a specific monitoring plan for the Microsoft Azure[©] platform.

The paper is structured as follows: Sect. 2 discusses the existing solutions. Section 3 presents the monitoring infrastructure, its components and artifacts. Section 4 presents a case study. Finally, Sect. 5 presents the conclusions and future work.

2 Related Work

Several studies whose aim has been to analyze the monitoring tools and approaches that are available (e.g., [5, 12]) and their weaknesses and needs have appeared over the last few years. Fatema et al. [12] report the results of a survey in which they analyze cloud and general purpose monitoring tools. They identify practical capabilities that an ideal monitoring tool should possess in order to fulfill the objectives of both cloud providers and customers in different cloud operational areas. They conclude that most general purpose monitoring tools were not designed with the cloud in mind, signifying that most monitoring capabilities (e.g. multi-tenancy, scalability, non-intrusiveness) are improved using cloud based monitoring tools. However, one of the drawbacks of cloud monitoring tools is their portability. This reinforces the fact that many cloud specific monitoring tools are commercial and vendor dependent, which makes the tools less flexible and portable and means that their results are neither extensible nor comparable to other platforms. Aceto et al. [5] analyze and discuss the properties of a monitoring system for the cloud. They conclude that cloud monitoring tools should have quality characteristics (e.g., scalability, elasticity, adaptability) that will enable them to tackle the challenges that cloud monitoring implies. However, they also conclude that current solutions still require considerable effort if desirable characteristics are to be attained.

Many cloud providers offer their customers the ability to monitor cloud services using monitoring tools available for CPU, storage and network [13]. These tools are

closely integrated with their own cloud solutions. They are only concerned with monitoring quality attributes for the hardware resources (CPU, storage, and network) and lack the ability to monitor application-specific QoS parameters and SLA requirements (i.e., latency, performance). In addition, the majority of commercial tools (e.g., CloudWatch, LogicMonitor) are not sufficiently flexible to allow service providers to extend the QoS parameters provided to monitor the fulfillment of SLAs.

Various approaches have also been proposed in academic environments. For instance, Emeakaroha et al. [14] propose an application monitoring architecture named Cloud Application SLA violation Detection architecture (CASViD). This architecture monitors and detects SLA violations on the application layer, and includes tools for resource allocation, scheduling, and deployment. Although their approach provides a good solution, it does not have a flexible means to change the NFRs and metrics to be monitored at runtime. Katsaros et al. [15] present a monitoring system that facilitates on-the-fly self-configuration in terms of both the monitoring time and the monitoring parameters. They propose the use of scripts to collect data; however, they do not specify how NFRs are matched with raw data gathered from scripts and how they interact with cloud services. Müller et al. [8] designed and implemented SALMonADA, a service-based system with which to monitor and analyze SLAs in order to provide an explanation of violations. They describe SLAs using a Monitoring Management Document (MMD) to be consumed by the monitoring infrastructure; however, the platform does not support those users who wish to choose alternative means to measure quality requirements. Smit et al. [16] present and implement an architecture using stream processing to provide service monitoring. They emphasize that their infrastructure is intended be used to monitor hybrid clouds and two tiered cloud architectures working on streaming data. The possibility of gathering information therefore depends on the information that can be provided by other solutions. Montes et al. [17] propose a cloud monitoring taxonomy, which is used as the basis to define a layered cloud monitoring architecture. They implement GMonE, a general-purpose cloud monitoring tool, which is claimed to cover all aspects of cloud monitoring by specifically addressing the needs of modern cloud infrastructures. Similarly, Povedano-Molina et al. [18] propose DARGOS, a distributed architecture for resource management and monitoring in clouds, which ensures an accurate measurement of physical and virtual resources in the cloud in an attempt to keep overheads down. However, the latter two approaches confront the provision of only physical and virtual resources and do not emphasize the specific quality aspects of SaaS. In summary, to the best of our knowledge commercial tools are mostly tightly coupled with certain cloud platforms, support the monitoring of specific NFRs, and have pre-established low-level metrics; they are therefore not sufficiently versatile to support the modification of NFRs or the customization of their operationalizations¹ at runtime. There are other proposals that

¹Operationalizing a measure consists of establishing a mapping between its generic description and the concepts represented in the software artifacts to be measured [30].

allow the verification of SLA compliance, but they are not enough flexible to support different operationalizations according to the specific cloud platform involved.

3 Monitoring Infrastructure

In this section, we present the *Monitoring Infrastructure* that has been designed to support the monitoring process defined in Cedillo et al. [11] (see Fig. 1). This infrastructure allows: (i) the specification and configuration of NFRs to be monitored; (ii) an interaction with cloud services to assess their quality at runtime; (iii) and the generation of reports containing any eventual SLA violations. In order to achieve these goals and provide the required degree of flexibility when defining NFR metrics, in addition to supporting different means to gather information from cloud services, we have defined a set of components and artifacts that conform to the monitoring infrastructure by using models@run.time. The solution is oriented to be applied in any platform, a detailed instantiation of the middleware to a defined platform is presented in Cedillo et al. [19].

The Monitoring Infrastructure has two main components: the *Monitoring Configurator* and the *Monitoring and Analysis Middleware*. The Monitoring Configurator uses the *Monitoring Requirements Model* and the *SaaS Quality Model* to configure the monitoring of services and obtain the *Runtime Quality Model*. The *Monitoring and Analysis Middleware* uses this Runtime Quality Model and relies on two engines: the *Measurements Engine*, which permits cloud service monitoring through the use of the raw service quality data gathered from cloud services and takes the measurements, and the *Analysis Engine*, which compares the expected values with the monitored values and can generate the SLA violations report. The details of each process and artifact are detailed in the following subsections.

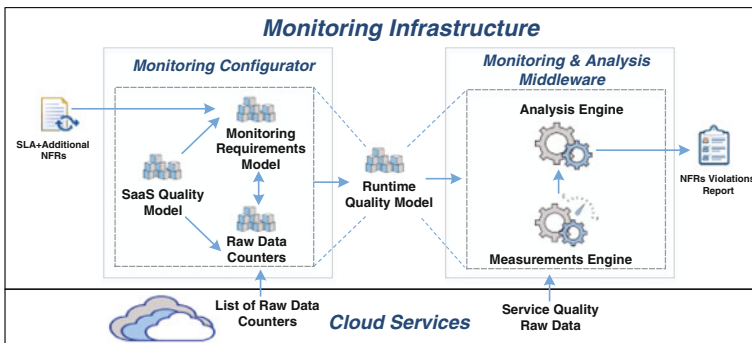


Fig. 1 Monitoring infrastructure

SLAParameters. Each SLAParameter can be measured by using metrics. The SLAParameter meta-class has an attribute named isSLATerm, which differentiates an SLA term from an NFR that is not included in the SLA. The Obligation meta-class contains two types of obligations: (i) a Service Level Objective, which is a guarantee of a particular state of SLA parameters in a given time period. (e.g. the average response time must be 5 ms) and (ii) The Action Guarantee, which specifies the provider’s commitment to doing something in a specific situation [21] (e.g. if a violation of a guarantee occurs, a notification is sent specifying a penalty). The values used as thresholds are obtained from the Action Guarantee meta-class (e.g. the response time must be <0.7 unless the transaction rate is >1000). In this meta-model, a metric can be measured by using the formula agreed by the parties. A more detailed specification of the WSLA used to define the meta-classes, with examples, can be found in Ludwig et al. [20].

SaaS Quality Model This model is aligned with the ISO/IEC 25010 standard (SQuaRE) [21]. Figure 3 shows the meta-model supporting the SaaS Quality Model. This model allows the definition of the whole set of Characteristics, Sub-characteristics, Attributes, their Impact (i.e., the relationships among attributes), and Metrics that specify how NFRs should be measured to assess the quality of cloud services. Each metric can be operationalized in different ways. A metric Operationalization can be considered at different Cloud Levels (i.e., SaaS, PaaS, IaaS). This is useful owing to the fact that there are a number of quality requirements (e.g., scalability, elasticity, security) that need to be monitored for different levels of service provision [5]. Moreover, it is important to specify the stakeholder that will use the monitoring information; for example, for a service provider, it may be interesting to know the average number of users requesting a service at a particular time. The purpose of having perspectives associated with each operationalization is to express whether a given operationalization is stakeholder-specific. This information is useful during the processes of comparing, improving measurements, or choosing different formulas with which to measure each NFR. The DirectMetricOperationalization meta-class represents a measure of an attribute that does not depend upon any other measure, whereas the

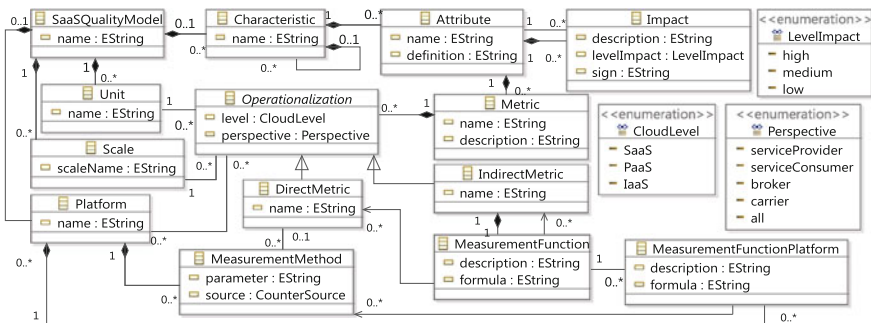


Fig. 3 SaaS quality meta-model

IndirectMetricOperationalization meta-class represents measures that are derived from other DirectMetricOperationalizations or IndirectMetricOperationalizations. The Platform and MeasurementMethod meta-classes have been added to the SaaS Quality Model to maintain a list of raw platform dependent data counters to facilitate the retrieval of information from a specific platform. Finally, the meta-model includes particularities of each operationalization, such as the Unit meta-class, which expresses the magnitude related to a particular quantity. The Scale meta-class represents a set of values with continuous or discrete properties used to map the operationalization.

Runtime Quality Model This is a model@run.time, which specifies the monitoring requirements, metrics, operationalizations, and configurations that will be used during the monitoring. Lehmann et al. [22] argue that meta-models at runtime must provide modeling constructs that will enable the definition of: (a) A prescriptive part of the model, specifying what the system should be like; (b) A descriptive part of the model specifying what the system is like; (c) Valid model modifications of the descriptive parts, executable at runtime; (d) Valid model modifications of the prescriptive parts, executable at runtime; (e) Causal connection, which is in the form of an information flow between the model and the entity being monitored. Figure 4 shows the Runtime Quality Meta-model, which is an extension of the SaaS Quality Model. It has many of the meta-classes included in the SaaS Quality Model described previously, plus meta-classes that represent the prescriptive part, the descriptive part, and the characteristics of the cloud platform that allow the causal connection.

The CloudService meta-class also describes the service to be monitored. The prescriptive part of the model thus includes the *Threshold*, which is obtained from the obligations part of the SLA, or an AdditionalNFR threshold set by the

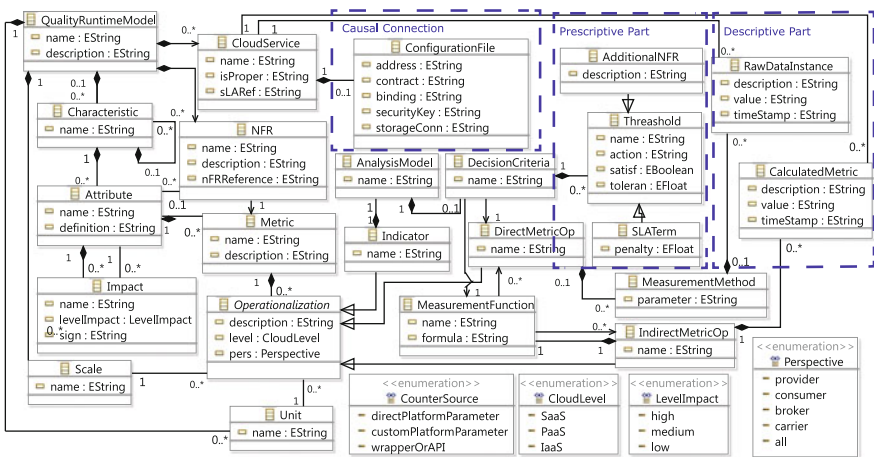


Fig. 4 Runtime quality meta-model

stakeholder. The descriptive part of the model is formed of the RawDataInstance meta-class, which contains the values captured directly from the cloud, and the CalculatedMetric meta-class, which contains the measurement results of the calculated metrics. The ConfigurationFile meta-class contains specific information for each platform that allows an interaction to take place between the monitoring infrastructure and the cloud service. It can therefore be considered as the class that is used to attain the causal connection between the monitoring infrastructure and services when a change needs to be reflected. Finally, the Indicator meta-class represents a measure that is derived from the other measures using an Analysis Model as a measurement approach [23]. In conclusion, the Runtime Quality Model allows our proposal to obtain the desirable characteristics related to flexibility and maintainability, since changes in the Runtime Quality Model can be easily reflected in the monitoring infrastructure.

Interaction Among Models Figure 5 shows the interactions among the models. The first interaction (1) occurs between the SaaS Quality Model and the Monitoring Requirements Model. Stakeholders can use the SaaS Quality Model, which contains a standardized classification of characteristics, sub-characteristics, metrics, and attributes, as support in order to define the Monitoring Requirements Model. The second interaction (2) then occurs between the Monitoring Requirements Model and the Runtime Quality Model. Here, the stakeholder uses the Monitoring Configurator Interface to capture the NFRs and metrics included in the Monitoring Requirements Model to define the Runtime Quality Model. Finally, the third interaction (3) occurs between the Runtime Quality Model and the SaaS Quality Model. This interaction allows the means used to gather information from cloud services to be specified. In this scenario, the SaaS Quality Model is useful as regards matching the high level attributes contained in the Monitoring Requirements Model with raw service quality data. Here, the SaaS Quality Model

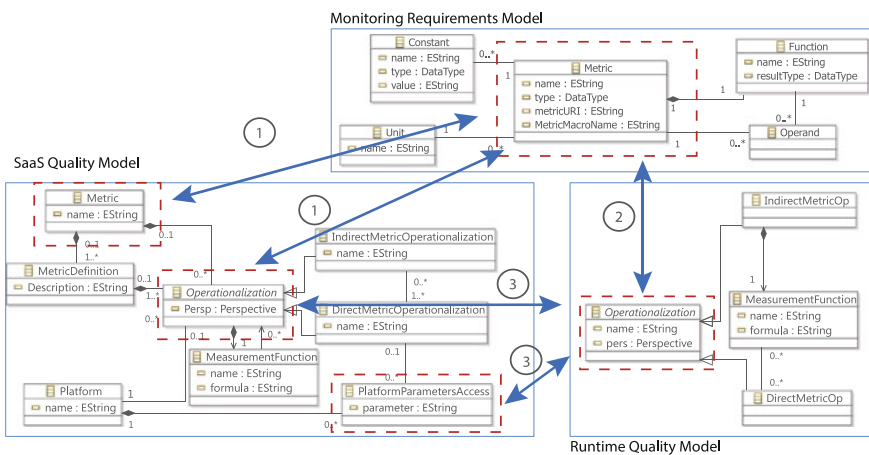


Fig. 5 Interaction among models

enables a choice to be made from among many equivalent operationalizations with different measurement methods, thus providing our approach with flexibility.

3.2 Monitoring and Analysis Middleware

The Monitoring and Analysis Middleware consists of the Measurements Engine, which uses the Runtime Quality Model obtained as result of the configuration as input, and this applies metrics with which to measure the quality of services. There is also the Analysis Engine, which permits the analysis of quality and reports SLA violations. A description of the Monitoring Middleware components has been addressed in Cedillo et al. [19].

4 Case Study

An exploratory case study was performed following the guidelines presented in Runeson et al. [24] in order to analyze the feasibility of the configuration task. The stages of the case study are: design, preparation, collection of data, and analysis of data, each of which is explained below. In this case study, we have used a metric to illustrate the configuration task. In Cedillo et al. [19], it can be found other examples to have a better idea about the configuration and application of other NFRs to this solution.

4.1 Design of the Case Study

The case study was designed by considering the five components proposed in Runeson et al. [24]: purpose of the study, underlying conceptual framework, research questions to be addressed, sampling strategy, and methods employed.

The purpose of this case study is to analyze the feasibility of configuring the monitoring of services by means of the Monitoring Configurator, and to use these configurations to generate the Runtime Quality Model. The Monitoring and Analysis Middleware will take this model as input to monitor the cloud services. The conceptual framework that links the phenomena to be studied is based on the Monitoring Process [11] and an infrastructure that supports this process (i.e., components, artifacts). The research questions to be addressed are: (a) is the strategy of configuring and matching the NFRs with quality raw data retrieved from cloud services to obtain the desired monitoring information useable and effective?; (b) what are the limitations of the monitoring configurator?

Here, the sampling strategy is based on monitoring configuration tests carried out by a subject who is an IT professional with programming skills and who has

been working as a Cloud Provider Service Specialist for two years. In accordance with Lethbridge et al. [25], we have applied the second degree of data collection techniques, in which the researcher directly collects raw data without interacting with the subject during the data collection.


In order to collect the monitoring information, we developed a prototype of the Monitoring and Analysis Middleware that allows the collection of raw data through the use of the Runtime Quality Model generated in the configuration task. The monitoring configuration was carried out as follows: the subject used the Monitoring Requirements Model to match NFRs with quality parameters and instructions that gather data from a service running in the cloud. The technique used to obtain feedback regarding the feasibility of the monitoring configuration performed was an analysis of the monitoring results obtained using a prototype of the Monitoring Engine to obtain the data needed to prove whether the values gathered were those expected by the subject.

4.2 Preparation of the Case Study

The context of this case study, was a test scenario in which the subject carried out the monitoring configuration. The SaaS Quality Model was used to support the matching between the NFRs to be monitored and the platform information. Once this information had been matched, it was possible to generate the Runtime Quality Model, which was then used by the Monitoring and Analysis Middleware to gather, measure and analyze quality data obtained from cloud services. The services used in this case study were implemented in compliance with an Open Reference Case (ORC) proposed in Ludwig et al. [20], which was used as an open source demonstrator to highlight the achievements of the European research project SLA@SOI. The ORC is an extension of the CoCoMe implementation [26], which provides a service oriented retail solution that can be used in a supermarket trading system to handle the sales and stocking process [27]. The set of services defined by ORC was deployed as a SaaS on the Microsoft Azure© platform. We considered the actions (i.e., create, read, update, and delete operations) related to the inventory service and the sales service. The objective was to configure the monitoring infrastructure in order to perform quality evaluations of cloud services. The NFRs to be monitored were reliability and latency.

Figure 6 presents an excerpt of an instance of the Monitoring Requirements. It shows the service, its operations (e.g.NewItemInventory) and the NFRs (SLAParameters). The NFRs to be monitored are the *reliability* and *latency* of the inventory and sales cloud services. *Reliability* is defined as “*the ability of an item to perform a required function under stated conditions for a stated time period*” [28]. Customers and suppliers often measure service reliability as Defective operations Per Million attempts (DPM) [29]. In this case study, the SLA term included the following clause: “the service could have a maximum of ten defective operations per million” (i.e., “99.999 % service reliability”). *Service latency* was, meanwhile,

Fig. 6 Monitoring requirements model instance excerpt



defined as “the time that has elapsed between a request and the corresponding response” [29], and thus “the maximum service latency is 130 ms”. The *Monitoring Requirements Model* includes the DPM metric which measures reliability. It was then necessary to select the DPM equivalent operationalization, which allows the measurement of the reliability NFR in cloud services deployed on the Microsoft Azure © platform. Our SaaS Quality Model contains three equivalent metric operationalizations (i.e. DPM1, DPM2, and DPM3). The subject select one of them depending on the Monitoring Requirements Model and the Raw Service Quality Data to be retrieved.

The operationalizations included in our SaaS Quality Model to calculate DPM are:

$$\text{DPM 1} = \frac{\text{Operations Attempted} - \text{Operations Successful}}{\text{Operations Attempted}} \times 10^6 \quad (1)$$

$$\text{DPM 2} = \frac{\text{Operations Failed}}{\text{Operations Attempted}} \times 10^6 \quad (2)$$

$$\text{DPM 2} = \frac{\text{Operations Failed}}{\text{Operations Successful} + \text{Operations Failed}} \times 10^6 \quad (3)$$

The subject can select an equivalent operationalization by considering the advantages and disadvantages of the selection (e.g. overheads). Once the Runtime Quality Model has been generated, the Monitoring and Analysis Middleware can collect information, measure data, and report SLA violations. Here, data is captured by using the Azure Diagnostics Service. However, this could change depending on the facilities of each cloud platform. Here, the subject can use one of the three equations (1)–(3) to match that selection with Diagnostics counters. Finally, the matched formula was used for the Monitoring and Analysis Middleware using Diagnostics counters.

4.3 Collection of Data

The data was collected in two stages: (1) when the subject carried out the configurations depending on the NFRs and matched these NFRs with raw platform-specific data counters to generate the Runtime Quality Model using the

PartitionKey	RowKey	Timestamp	Value	Name	Service
63560457099254...	63560457099254...	25/02/2015 9:31...	0	Downtime	CoCoMe-Sales
63560457099971...	63560457099971...	25/02/2015 9:33...	61855	Defective Opera...	CoCoMe-Sales
63560457218420...	63560457218420...	25/02/2015 9:33...	93,81	Reliability	CoCoMe-Sales
63560457219434...	63560457219434...	25/02/2015 9:33...	0	Latency	CoCoMe-Inventory
63560457220186...	63560457220186...	25/02/2015 9:33...	0	Downtime	CoCoMe-Inventory
63560457337688...	63560457337688...	25/02/2015 9:35...	69405	Defective Opera...	CoCoMe-Inventory
63560457338588...	63560457338588...	25/02/2015 9:35...	93,0595	Reliability	CoCoMe-Inventory

Fig. 7 Metrics calculated by using the monitoring infrastructure

SaaS Quality Model; (2) when the monitoring engine gathered and measured information provided by cloud services based on the Runtime Quality Model. A prototype of the Monitoring and Analysis Middleware was implemented as a Microsoft Azure cloud service, which stores the results in a data base. Figure 7 shows the results with the calculated metrics.

4.4 Analysis of Data

The monitoring configuration was analyzed so as to address our research questions. The subject used the Monitoring Requirements Model, which contained the NFRs to be monitored, and their metrics and thresholds. The subject then matched the metrics with the appropriate operationalizations specific to the platform. In order to illustrate the process used to monitor the reliability, the other NFRs were monitored following analogous steps. The reliability threshold was 99.999 %, and we the considered operationalization (1) which was set up by matching formula (4) with the Azure Counters:

- RequestsTotal=@“\ASP.NET Applications(_Total_)Requests Total”
- OperationsSuccessful=@“\ASP.NET Applications(_Total_)Requests Succeeded”

$$DPM = \frac{RequestsTotal - RequestsSucceeded}{RequestsTotal} \tag{4}$$

The Runtime Quality Model should include Formula (4). When checking whether the monitoring infrastructure would be able to monitor the behavior of the cloud services by using the runtime quality model, we intentionally introduced exceptions into the ORC services’ source code to generate reliability and latency problems.

It was necessary to determine whether the configuration gathers the expected information from the cloud services by using the Runtime Quality Model and to find possible limitations or inaccurate results. Here, we have concluded that the

Runtime Quality Model produced the expected values shown in the table presented in Fig. 7, in which the exceptions introduced were reflected in the monitoring results (the reliability offered was 99.999 % and the actual Reliability was 93.0595 % for the inventory service, signifying that the SLA was violated).

4.5 Case Study Conclusions and Lessons Learned

With regard to the first research question stated for this case study, we provide support to help the configuration of NFRs to be monitored using our approach and that the configuration was effective as regards monitoring Azure cloud services. Moreover, the suitability of this approach is shown by the fact that it is feasible to use the Monitoring Configurator to match the NFRs included in the Monitoring Requirements Model with the raw service quality data gathered from the cloud service and provide the expected information. With regard to the second research question, the Monitoring Infrastructure is able to detect SLA violations from a wide range of NFRs. However, it is important to take into account that not all the NFRs can be monitored owing to the restriction of the infrastructure that provides the raw service quality data from the services. One solution to this issue would be to use wrappers for services in order to capture the information required in a customized manner, which constitutes one of the next steps in our research.

As lessons learned this case study has allowed us to observe the potentialities and limitations of our proposal. The monitoring configurator allows a wide variety of operationalizations and platform counters to be matched. However, it depends on the facilities used to provide raw service quality data. During the execution of the case study, several aspects related to how the configuration can be facilitated have been discovered. For example, the SaaS Quality Model provides a simple means to choose the operationalizations and it is possible to add operationalizations to the SaaS Quality Model, which represents a knowledge base that saves efforts and minimizes possible mistakes when the configuring task is being carried out.

5 Conclusions and Future Work

In this paper, we have presented a monitoring infrastructure for cloud services, which allows data to be retrieved from cloud services in order to calculate monitoring metrics and eventually report non-compliance with the SLA. The monitoring infrastructure uses the Runtime Quality Model, which is generated by using two additional models: the Monitoring Requirements Model and the SaaS Quality Model. The feasibility of the approach has been illustrated by means of a case study which shows the monitoring of services deployed on the Azure platform.

The use of models@run.time provides flexibility and eases maintainability when the SLA and additional NFRs to be monitored change. Moreover, the facility of

changing the model and not the monitoring infrastructure makes it easier to operate and understand when they are not familiar with the middleware implementation.

As future work, we plan to deliver our Monitoring and Analysis Middleware in other platforms (e.g. Amazon AWS, Google) to be able to monitor and analyze services deployed in these platforms. We also plan to carry out a systematic review of the quality characteristics, sub-characteristics, attributes, and metrics of cloud services. The findings will be included in the SaaS Quality Model in order to study the monitoring mechanisms provided by other commonly used cloud platforms such as Google App Engine or Amazon AWS. Moreover, we plan to study generic means to encapsulate the raw data collected from the cloud services in order to obtain common interfaces for many platforms (e.g., APIs, proxies, plugins). Finally, we plan to improve the efficiency of the proposal by taking in account issues such as overheads, security, etc. and to empirically validate the approach using controlled experiments.

Acknowledgments This research has been supported by the Value@Cloud project (TIN2013-46300-R), Scholarship Program Senescyt-Ecuador, NSERC (Natural Sciences and Engineering Research Council of Canada) and Microsoft Azure for Research Award Program.

References

1. Sriram, I., Khajeh-Hosseini, A.: Research agenda in cloud technologies. In: 1st ACM Symposium on Cloud Computing, SOCC, pp. 1–11 (2010)
2. Song, J., Han, F., Yan, Z., Liu, G., Zhu, Z.: A SaaSify tool for converting traditional web-based apps to SaaS application. In: 4th International Conference on Cloud Computing, CLOUD, pp. 396–403 (2011)
3. Baset, S.A.: Cloud SLAs: present and future. *ACM SIGOPS Oper. Syst.* **46**, 57–66 (2012)
4. Hassan, M., Song, B., Huh, E.-N.: A market-oriented dynamic collaborative cloud services platform. *Ann. Telecommun.* **65**, 669–688 (2010)
5. Aceto, G., Botta, A., de Donato, W., Pescapè, A.: Cloud monitoring: a survey. *Comput. Netw.* **57**, 2093–2115 (2013)
6. Shao, J., Wei, H., Wang, Q., Mei, H.: A runtime model based monitoring approach for cloud. In: International Conference on Cloud Computing (CLOUD), pp. 313–320 (2010)
7. Foster, I., Zhao, Y., Raicu, I., Lu, S.: Cloud computing and grid computing 360-degree compared. In: Grid Computing Environments Workshop (GCE 08), pp. 1–10 (2008)
8. Muller, C., Oriol, M., Franch, X., Marco, J., Resinas, M., Ruiz-Cortés, A., Rodríguez, M.: Comprehensive explanation of SLA violations at runtime. *IEEE Trans. Serv. Comput.* **7**, 168–183 (2014)
9. Baresi, L., Ghezzi, C.: The disappearing boundary between development-time and run-time. In: Workshop on the Future of Software Engineering Research FSE/SDP, pp. 17–22. ACM, USA (2010)
10. Giese, H., Bencomo, N., Pasquale, L., Ramirez, A., Inverardi, P., Wätzoldt, S., Clarke, S.: Living with Uncertainty in the Age of Runtime Models. http://dx.doi.org/10.1007/978-3-319-08915-7_3
11. Cedillo, P., Gonzalez-Huerta, J., Insfrán, E., Abrahao, S.: Towards monitoring cloud services using Models@run.time. In: Workshop on Models@run.time, MODELS, pp. 31–40, Spain (2014)

12. Fatema, K., Emeakaroha, V.C., Healy, P.D., Morrison, J.P., Lynn, T.: A survey of cloud monitoring tools: taxonomy, capabilities and objectives (2014)
13. Alhamazani, K., Ranjan, R., Mitra, K., Rabhi, F., Jayaraman, P.P., Khan, S.U., Guabtni, A., Bhatnagar, V.: An overview of the commercial cloud monitoring tools: research dimensions, design issues, and state-of-the-art. *Computing*, pp. 1–21 (2014)
14. Emeakaroha, V.C., Ferreto, T.C., Netto, M.A.S., Brandic, I., De Rose, C.A.F.: CASViD: application level monitoring for SLA violation detection in clouds. In: *Computer Software and Applications Conference (COMPSAC)*, pp. 499–508 (2012)
15. Katsaros, G., Kousiouris, G., Gogouvtis, S.V., Kyriazis, D., Menychtas, A., Varvarigou, T.: A self-adaptive hierarchical monitoring mechanism for Clouds. *J. Syst. Softw.* **85**, 1029–1041 (2012)
16. Smit, M., Simmons, B., Litoiu, M.: Distributed, application-level monitoring for heterogeneous clouds using stream processing. *Future Gener. Comput. Syst.* **29**, 2103–2114 (2013)
17. Montes, J., Sánchez, A., Memishi, B., Pérez, M.S., Antoniu, G.: GMonE: a complete approach to cloud monitoring. *Future Gener. Comput. Syst.* **29**, 2026–2040 (2013)
18. Povedano-Molina, J., Lopez-Vega, J.M., Lopez-Soler, J.M., Corradi, A., Foschini, L.: DARGOS: a highly adaptable and scalable monitoring architecture for multi-tenant Clouds. *Future Gener. Comput. Syst.* **29**, 2041–2056 (2013)
19. Cedillo, P., Jimenez-Gomez, J., Abrahao, S., Insfran, E.: Towards a monitoring middleware for Cloud services. In: *International Conference on Services Computing (SCC)*, NY, USA (2015)
20. Ludwig, H., Keller, A.: *Web Service Level Agreement (WSLA) Language Specification*, pp. 1–110 (2003)
21. ISO/IEC: *ISO/IEC 25010 Systems and Software Quality Requirements and Evaluation (SQuaRE)—System and software quality models* (2011)
22. Lehmann, G., Blumendorf, M., Trollmann, F., Albayrak, S.: *Meta-modeling runtime models*. In: *International Conference on Models in Software Engineering*, pp. 209–223. Springer, Berlin (2010)
23. García, F., Bertoa, M.F., Calero, C., Vallecillo, A., Ruíz, F., Piattini, M., Genero, M.: *Towards a consistent terminology for software measurement* (2006)
24. Runeson, P., Höst, M.: *Guidelines for conducting and reporting case study research in software engineering*. *Empirical Softw. Eng.* **14**, 131–164 (2009)
25. Lethbridge, T.C., Sim, S.E., Singer, J.: *Studying software engineers: data collection techniques for software field studies*. *Empirical Softw. Eng.* **10**, 311–341 (2005)
26. Herold, S., Klus, H., Welsch, Y., Deiters, C., Rausch, A., Reussner, R., Krogmann, K., Koziolk, H., Mirandola, R., Hummel, B., Meisinger, M., Pfaller, C.: *CoCoMe—the common component modeling example*. Presented at the (2008)
27. Wieder, P., Butler, J.M., Theilmann, W., Yahyapour, R. (eds.): *SLAs for Cloud Computing*. Springer, New York (2011)
28. *Quality Excellence for Suppliers of Telecommunications Forum (Quest Forum)*, TL 9000 Quality Management System Measurements Handbook 5.0 (2012)
29. Bauer, E., Adams, R.: *Service Quality of Cloud-Based Applications*. Wiley, Hoboken (2013)
30. Fernandez, A., Abrahão, S., Insfran, E.: *A web usability evaluation process for model-driven web development*. In: *International Conference on Advanced Information Systems Engineering*, pp. 108–122 (2011)