



Identification of Comparison Key Elements and Their Relationships for Cloud Service Selection

Anis Ahmed Nacer^(✉), Olivier Perrin, and François Charoy

Université de Lorraine, CNRS, Inria, LORIA, 54000 Nancy, France
anis.ahmed-nacer@inria.fr, {olivier.perrin, francois.charoy}@loria.fr

Abstract. Nowadays, the cloud computing industry is enjoying an exponential growth, where several cloud service providers compete to be one of the market leaders. Usually, providers offering similar services use different non-functional attributes to describe them. Thus, given the heterogeneity and diversity of services descriptions, the selection process of the appropriate cloud service becomes challenging. Architects no longer know what criteria to use to make the suitable cloud services selection. In this paper, we highlight the challenge of identifying key elements of comparisons and their relationship for selecting cloud services. Further, we propose a methodology to solve this issue based on real data available from service providers and benchmark work. Our methodology is validated based on two case studies of cloud relational databases and cloud queuing services.

Keywords: Cloud computing · Cloud service selection · Microservices

1 Introduction

Nowadays, cloud service providers present their services as Add-ons plans that are described according to a set of Non-Functional Attributes (NFAs). NFAs are settings that describe how a service is configured. They can influence Non-Functional Requirements (NFRs) such as performance, security and availability. The competition between the providers and the lack of standard have led them to describe their services differently to attract consumers. Service plans descriptions are different, incomplete, and use different designation for the same NFA. Further, these descriptions do not specify the relationship or impact between the comparisons' key elements (NFAs and NFRs). Thus, Choosing the most appropriate service to meet the NFR of architects becomes a challenging task. The selection of the most appropriate cloud service depends heavily on the choice of comparisons' key elements and their relationships. The comparisons' key elements of service plans are a set of NRFs and NFAs. In practice, incorrect selection of comparisons' key elements and their relationships results in an incorrect service evaluation. Therefore, comparisons' key elements and their relationships must be carefully described at the outset of the assessment.

The main challenge when selecting cloud services is to determine what are NFRs and NFAs that should be considered and how NFAs influence NFRs. In this paper, we propose a method that relies on architect input and service analysis. It consists of the following steps: (1) Understand the architects' requirements and how they reason regarding the choice of service plans for their applications. To do so, we propose to conduct semi-structured interviews with architects; (2) Identify the key elements of the comparison that meet the architects' requirements and the relationship between them. For that we propose to review service provider plans, works on cloud service benchmarks and literature reviews (3) Ensure the completeness of the list of key elements for comparison and their relationship. To do so we propose to conduct an empirical study with the architects. The outcomes of our methodology are the identification of a set of NFAs, NFRs and their relationships that we can use to evaluate cloud services.

The paper is organized as follows: Sect. 2 describes the proposed methodology for selecting the comparisons' key elements and their relationships to the selection of cloud services. Section 3 presents the validation of our methodology Sects. 4 and 5 present respectively the related work and the conclusion.

2 Proposed Method

An important step in selecting cloud services is to identify NFR, NFA and their relationships. Figure 1 gives an overview of the proposed methodology.

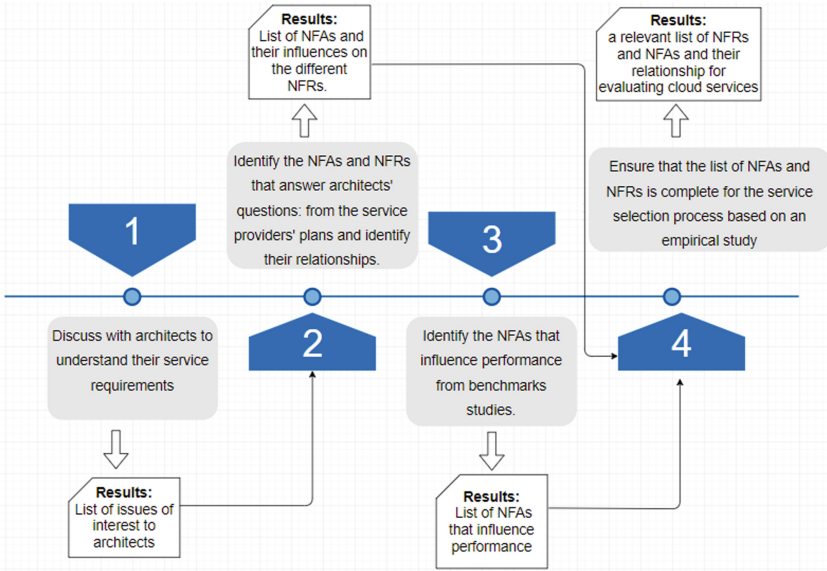


Fig. 1. Methodology for identifying NFAs and NFRs and their relationships for cloud service selection problems

2.1 Identification of Architects' Requirements

As a primary step, we organized interviews with some architects to clarify their minimum requirements and expectations from the different services. We asked them about their overall service requirements through open-ended questions like “what questions should they answer to choose the best plans from cloud service providers for their applications?” From these interviews, we have compiled a list of questions that architects consider when choosing cloud services, including the following as examples: how to predict capacity? Are performance specifications aligned with architects' expectations for cloud services? Do service plans provide replication and fail-over support?

The aim of this study is twofold: (1) group common issues for the generic evaluation of different services offered by different providers; and (2) establish the NFR that are used to assess services. The questions are grouped according to their influence on these requirements.

2.2 Identification of Attributes from Service Plans

As a second step, we need to identify the attributes that help to evaluate and match cloudware service plans to architect's requirements. Depending on plans, service provider indicates the presence or absence of some operational configuration. Evaluating plans on an absolute scale remains difficult. We identified attributes associated to service plans from providers that give access to their services on Heroku, IBM Cloud and Azure. We classified them into 4 categories, which are as follows: **(1) Capacity attributes** that provide information about the performance and the capacity of the service. We can use them to calibrate the plan according to the requirements of the application; **(2) Functional attributes** that give information about the functional coverage of the service. A database service for instance may include supplementary functions, including monitoring of certain features or auditing of functions; **(3) Service attributes** provide information on service level agreements (SLAs) such as SLA on support availability, SLA on service availability and protocols supported by the service; **(4) Technical attributes** include the attributes that describe how the service is deployed or operated. They have different impact on NFR of the service. To determine the impact of NFAs on NFRs, we first discussed with architects about their experiences with service deployment. Also, we reviewed the technical documentation of cloud services to understand why a NFA was advertised by cloud providers. By this way, we understand its influence on the NFR of the service. Determining the rate of influence of NFAs on NFRs is difficult. It depends on factors whose information is not available from the service providers. This includes the deployment of the provider's architecture: the number of nodes on which the service is deployed and the technologies used to develop the service. In addition, the influence of NFAs on NFRs depends on the nature of the applications for which the service is intended. Due to these difficulties, we only indicate the positive, neutral or negative impact on the technical attribute (+, =, -). Table 1 gives an example of a database service's technical attributes and their impact on NFRs.

Table 1. Example of the influence of certain NFAs on NFRs for relational cloud database services

Attributes	Performance	Availability	Reliability	Security	Scalability
RAM	+	=	=	=	=
IOPS	+	=	=	=	=
Backup	-	+	+	=	=
Rollback	=	+	=	=	=
Replication	-	+	+	=	=

2.3 Identification of Attributes Based on Benchmark Works

All NFRs, except performance, can be assessed from the NFAs supplied by cloud service providers. Performance should be considered as a particular NFR. The only way to identify the attributes that have an impact on this NFR, is to refer to benchmarks. In order to understand the attributes that influence service performance, we reviewed previous benchmarking work¹ that compares the performance of different cloud services for different use cases, under different constraints and experience configurations. This study is interesting but incomplete because most benchmarking works do not provide all the attributes that influence performance in their experiments. In experiments for the same use case, with the same constraints and experimental configurations, we observed that the performance results are different. We do not know if this is due to external factors undeclared attributes or technical inaccuracies. To complete our study, we reviewed previous work that focused on a structured analysis of the fundamental principles of variation and predictability of service provider performance.

2.4 Selection of NFAs and NFRs from an Empirical Study

To ensure the completeness and the relevance of the collected technical attributes, we propose to conduct an empirical study. The aim is to understand how architects decide the types of components to develop a new application, and the criteria that guide their choices. We are particularly interested in core components (database, messaging, caching, indexing, monitoring...) that can be deployed as servers or consumed as services in the cloud. We conducted a semi structured interview, as done in [1], on a general perception of services with seven software architects. All of them have more than four years of experience in cloud computing environments and have participated in several projects covering different application areas such as data engineering, software/systems development, engineering systems and test infrastructure. Individual interviews were performed. Each lasted around one hour. The interviews are recorded, transcribed and synthesized. The architects' responses are assessed on the basis of their technical knowledge and experience. The survey for this study is built on

¹ <https://www.2ndwatch.com/blog/benchmarking-amazon-aurora/>.

37 questions divided into three parts: questions on the architect’s experience, questions on the overall application and questions on each service/component². For example, when asked: “Why do you want to move from IBM’s provider database service to Amazon for your application?” Architects A1, A2, A3, A4 and A6 replied: “IBM does not provide Virtual Private Cloud (VPC) for its services and this motivates us to switch to Amazon services” and architects A5 and A7 replied: “We consider VPC important but not important enough to migrate our service to another cloud service provider”. These responses conclude that the NFA “VPC” is necessary for the architects’ application. Therefore, VPC can be an important NFA for cloud service evaluation according to the architects’ application requirements. We asked the same types of questions about service selection and collected the NFA that were missing for service evaluation, ensuring at the same time that the list of NFA we had collected was complete for service evaluation. To summarize, we filtered the NFA that were collected in the previous steps and retained only the ones that are relevant to the evaluation of services. In addition, we identify missing attributes that could not be collected in the previous steps. Using our methodology, we aim to identify a relevant list of NFAs, NFRs and their relationships for which cloud services will be evaluated.

3 Empirical Validation of the Method

To validate the approach described in Sect. 2, we conducted two case studies: SQL database services and queue services. We examine these services on the basis of actual configuration data available from service providers and benchmarking works. Due to space limitation, only the first use case (SQL database) is presented. The second use case is shared here³. In the rest of this section, we apply our methodology to identify the key elements of the comparison and their relationships to the SQL cloud service.

– **Identification of Architects’ Requirements:** For this step, we focus on questions raised in Subsect. 2.1. The questions are grouped according to their influence on the two NFRs reliability and availability. Regarding availability, architects raised a set of questions such as: What are the regions available for service? Are they adequate in terms of legal and regulatory requirements? Do service plans provide replication and fail over support? For reliability, the architects raised a set of questions such as: What is the frequency/severity of failures for the service? What is the mean time between failures (MTBF)? What is the mean time to repair the system (MTTR)? Note that the findings of this study are generic and can be reused for all types of services. The full results of this study are described in⁴.

² <https://docs.google.com/document/d/11LTsFJCTSqNIX5DJPYgdU9NRW8IFyd6eGiRimiY2JeQ/view>.

³ https://drive.google.com/file/d/1tBouKaagH8IMJdoYDPDri7U7Oo_iHUSY/view?ths=true.

⁴ <https://drive.google.com/file/d/1s7xy3u-voLkO9XKk03K5s2yqdHgaHYfFi/view?usp=sharingforacompleteanddetailedlistofthequestionsofthisstudy>.

– **Result of Identification of NFAs and their influence on NFRs:**

Many SQL database service solutions are available and are deployed by different cloud service providers. The interesting attributes of these services are the following: **(1) The capacity attributes:** the only ones that can be identified are: **(i)** Storage capacity in Gb; **(ii)** IOPS that corresponds to the underlying disk performance; **(iii)** The maximum simultaneous connection to the database; and **(iv)** The row limit. **(2) The functional attributes:** include the monitoring and the audit. **(3) The technical attributes:** are more diverse and heterogeneous and they may have different impact on NFA. As explained previously, with the exception of performance, for which additional benchmarking may be required, there are only two ways to determine the impact of NFAs on cloud services NFRs: use the expertise of the architects who deployed the service or consult the technical documentation^{5,6,7}. Given the difficulty of determining the level of influence of the NFA on NFRs (explained in Sect. 2.2), we present in Table 2 the technical attributes with the estimated positive, neutral or negative impact of NFR (+, =, -). At this stage this is all we can say about this impact.

Table 2. Influences of NFAs on NFRs for relational cloud database services

Attributes	Performance	Availability	Reliability	Security	Scalability
RAM	+	=	=	=	=
IOPS	+	=	=	=	=
SSD	+	=	=	=	=
Server cores	+	=	=	=	=
Backup	-	+	+	=	=
Rollback	=	+	=	=	=
Replication	-	+	+	=	=
Single tenant	+	=	=	+	=
Multi tenants	-	=	-	+	=
Encryption	=	=	=	+	=
VPC	=	=	=	+	=
Multi region	=	+	+	=	=
Scaling auto increase	=	=	=	=	+
Scaling not supported	=	=	=	=	-
Scaling converted to other storage types	=	=	=	=	+
Read replicas	-	=	=	=	+

⁵ <https://docs.oracle.com/cloud/latest/mysql-cloud/UOMCS/UOMCS.pdf/>.

⁶ <https://cloud.google.com/sql/docs/mysql/>.

⁷ <https://medium.com/@lakshmanLD/comparison-of-mysql-across-aws-azure-and-gcp-19af2d208d9a>.

– **Result of identification of attributes based on benchmarks:** Cloud SQL services run on providers’ virtual machines. These virtual machines have a direct impact on the SQL services that run on them. Therefore, NFAs that influence the performance of virtual machines also influence the performance of SQL services. So we reviewed previous work [7] that conducted a large-scale literature review to collect and codify existing research on the predictability of public IaaS cloud performance. We identified the following attributes are relevant to assess the performance of cloud SQL services: **(1)** the CPU model, the hardware heterogeneity and the tenancy model; **(2)** temporal and geographic (region) factors; **(3)** the number of nodes and the number of containers on which the service is deployed. To complete our study, we reviewed previous benchmarking work that has been done directly on cloud SQL services [2, 6]. We collected the following attributes in addition to those we previously identified on the performance of cloud SQL services: **(1)** user connections are a resource that can limit performance; **(2)** replication and failover significantly reduce service performance; **(3)** the location of the data center influences performance; **(4)** number of read replicas improves performance of SQL cloud database services.

– **Result of selection of NFAs and NFRs from an empirical study**
 We carried out an empirical study with the architects on cloud SQL database services. A set of responses are collected. For example, to the question: “Why did you choose the High Availability (HA) option for your service when it costs about twice as much as a regular instance? architects A1, A2, A5 and A6 replied: “The MySQL service is a business service in our application. It should be very available. This is the critical part of the application. A failure of this service will cause the failure of the entire application. A failure of this service will cause the failure of the entire application. A fail over is necessary to maintain this service available when the instance is “blocked”. It has happened that the primary instance is “blocked” and that a restart of the instance takes up to 30 min (supplier ticket). To avoid these downtimes, we have opted for the HA option” and architects A4, A3 and A7 replied: “A complete failure of a zone is probably very rare, so HA is an important option, but not so important as to be indispensable.” We asked several questions about the selection of cloud SQL database services and ensured that the list of NFAs previously collected was relevant to our study. Based on previous studies and the architects answers, a relevant list of NFAs and their influences on NFRs was developed (as shown in Table 2) to select SQL cloud services based on the application requirements of the architects.

4 Related Work

Several approaches have been proposed to select the best service provider plans. These approaches use NFAs to evaluate cloud services. In this section, we examine the NFAs used in these approaches. We identified two main approaches, non-benchmark-based and benchmark-based. On one hand the Non-benchmark-based approaches [5, 9, 10] use different attributes such as: availability rate, execution time, instance starting time and mean time between failures (MTBF).

These attributes are difficult to express and can be addressed in several ways. Thus using them to solve the cloud service selection problem is of little relevance. On the other hand, several benchmark-based studies have been conducted. In [4], Smart CloudBench, a generic benchmarking tool allows to compare the offers available on the IaaS cloud market and monitor their performance. However, it is limited to IaaS and does not consider more heterogeneous middleware services. In [8], CloudCmp, a benchmarking tool is presented for cost comparison and performance measures of many services between different providers. In [3], a benchmarking methodology is applied to only four cloud storage service offers for specific workloads. The shortcoming of these approaches lies in a consideration of performance and price only, without considering other NFRs such as security, availability, reliability or security.

5 Conclusion

In this paper, we proposed a methodology for identifying key elements of comparisons and their relationship for cloud service selection. It is based on actual data available from service providers and benchmarks. First, we conducted a survey among architects to understand their requirements for selecting cloud services. Further, we identified issues that interest architects and we ranked them according to their influence on NFRs. Then, we identified the attributes of the service provider plans that answer the architects' questions and their influence on the NFRs from the technical documentation. Moreover, we reviewed the work on cloud benchmarking to collect the attributes that influence performance. Finally, we conducted an empirical study to ensure that the list of attributes we identified was complete and to add those that were missing. As a future work, we will use the results of this study to assess cloud services and cloud service composition by application type.

References

1. Ameller, D., Ayala, C., Cabot, J., Franch, X.: Non-functional requirements in architectural decision making. *IEEE Softw.* **30**(2), 61–67 (2012)
2. Bernstein, P.A., et al.: Adapting microsoft SQL server for cloud computing. In: *IEEE 27th International Conference on Data Engineering*, pp. 1255–1263. IEEE (2011)
3. Bocchi, E., Mellia, M., Sarni, S.: Cloud storage service benchmarking: methodologies and experimentations. In: *IEEE 3rd International Conference on Cloud Networking (CloudNet)*, pp. 395–400. IEEE (2014)
4. Chhetri, M.B., Chichin, S., Vo, Q.B., Kowalczyk, R.: Smart CloudBench—a framework for evaluating cloud infrastructure performance. *Inf. Syst. Front.* **18**(3), 413–428 (2016)
5. Karim, R., Ding, C., Miri, A.: An end-to-end QoS mapping approach for cloud service selection. In: *IEEE Ninth World Congress on Services (SERVICES)*, pp. 341–348. IEEE (2013)

6. Lang, W., Bertsch, F., DeWitt, D.J., Ellis, N.: Microsoft Azure SQL database telemetry. In: Proceedings of the Sixth ACM Symposium on Cloud Computing, pp. 189–194. ACM (2015)
7. Leitner, P., Cito, J.: Patterns in the chaos—a study of performance variation and predictability in public IaaS clouds. *ACM Trans. Internet Technol. (TOIT)* **16**(3), 15 (2016)
8. Li, A., Yang, X., Kandula, S., Zhang, M.: CloudCmp: comparing public cloud providers. In: Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement, pp. 1–14. ACM (2010)
9. Wagle, S.S., Guzek, M., Bouvry, P., Bisdorff, R.: An evaluation model for selecting cloud services from commercially available cloud providers. In: IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom), pp. 107–114. IEEE (2015)
10. Zheng, X., Martin, P., Brohman, K., Da Xu, L.: CLOUDQUAL: a quality model for cloud services. *IEEE Trans. Ind. Inf.* **10**(2), 1527–1536 (2014)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

