The Role of Experimental Exploration in Cloud Migration for SMEs

Frank Fowley[1], Divyaa Manimaran Elango[1], Hany Magar[1] and Claus Pahl[2]
[1]*IC4, Dublin City University, Dublin 9, Ireland*
[2]*Free University of Bozen-Bolzano, Bolzano, Italy*

Keywords: Cloud Migration, Experiment, Prototyping, Migration Method, SME, Cloud Architecture, Cloud Cost Model.

Abstract: The migration of IT systems to the cloud is still a problem, in particular for smaller companies without much cloud expertise. Generally, some expected benefits are defined and an awareness of potential problems does exist to some extent in the organisations. However, this is often not sufficient to confidently embark on a full migration project in the cloud. While discussions and conceptual analyses can help to some extent, we explore here the suitability of feasibility studies with experimental explorations at the core. These studies would typically cost 5% of the overall migration cost based on our use cases, but can help with a reliable risk assessment. It can clarify how much of the expectations and intentions can materialise in the cloud. The cost of the migration, but, more importantly, the cost of operating an IT system in the cloud can be estimated. Using a feasibility study with an experimental core based on a partial prototype delivers much more reliable figures regarding configurations, quality-of-service and costing than a theoretical analysis could deliver.

1 INTRODUCTION

While cloud computing is established, the migration of systems to the cloud (Jamshidi et al., 2013) is still a problem for small and medium-sized enterprises (SMEs) without cloud expertise (Giardino et al., 2015). These need to rely on support from consultants and solution providers. Expected benefits are known and an awareness of potential problems exists. However, this is often not sufficient to confidently embark on a full migration, which requires estimates of the migration costs as well as costs of operating software in the cloud within given quality limits.

Consultants and solution providers offer discussions and analyses that can help to some extent. However, some assumptions rely on expert knowledge from similar cases, but might not always be fully reliable. We explore here the suitability of feasibility studies with some experimental exploration at the core. These studies are a worthwhile investment, and would typically be 5% of the overall migration cost. A second application context where this applies is migration to another cloud architecture setting. So, rather than cloud-onboarding, the source architecture is already cloud-based. An example is an IaaS to PaaS migration from a basic, virtualized on-premise system into a fully cloud-native architecture. A proper project scoping for the migration is a key problem (Son, 2013), because of misconceptions and unclear technical expectations. Migration frameworks and case studies have been reported in the literature by academics and practitioners from industry (Jamshidi et al., 2013; Gholami et al., 2016; Jamshidi et al., 2016), but how to reliably estimate a proper 'right-scaling' of a cloud deployment remains unclear.

The benefit of feasibility studies is a reliable risk assessment. It can clarify how much of the expectations and intentions can materialise. The cost of the migration (Arshad et al., 2015), but, more importantly, the cost of operating an IT system in the cloud can be estimated (Jamshidi et al., 2014). It also helps to better understand technical cloud architecture concerns. Another question is a re-engineering one. What is migratable and what is the extent of refactoring necessary to make migration work? Often, re-engineering software in order to modernise and adapt to cloud constraints is needed. Using a feasibility study with an experimental core based on a partial prototype of the proposed cloud software system (Li et al., 2011) delivers much more reliable figures regarding configurations, quality-of-service and costing than a theoretical analysis could delivery.

We use case studies that we carried out as independent consultants with SMEs in the software sector using this model to validate the benefits.

The paper is organized as follows. In Section 2, we introduce a migration feasibility framework. Then, we define the aims of the feasibility studies. Practical concerns are presented in Section 4. In Section 5, we discuss observations.

Fowley, F., Elango, D., Magar, H. and Pahl, C.
The Role of Experimental Exploration in Cloud Migration for SMEs.
DOI: 10.5220/0006234503290334
In *Proceedings of the 7th International Conference on Cloud Computing and Services Science (CLOSER 2017)*, pages 301-306
ISBN: 978-989-758-243-1

2 MIGRATION FEASIBILITY FRAMEWORK

For the migration feasibility studies, we used a pattern-based approach to define a draft migration plan. This involves (i) the source architecture of the system, (ii) target architecture options in the cloud, and (iii) high-level architectural transformations for the different target architectures (Xiong et al., 2013). We use a catalogue of migration patterns to describe simple architectural transformations for specific scenarios (e.g. for simple cloudification in an IaaS solution). This defines a staged process based on a migration path which in the individual steps are driven by selection criteria (e.g., time to market or introduction of new capabilities). A sample pattern is Multi-Cloud Relocation (Jamshidi et al., 2014), see Fig. 1.

- Definition: A component re-hosted (or relocated) on a cloud platform is enhanced by using the environmental services of the other cloud platforms.
- Problem: Availability of an application needs to be enhanced without architecture change, and without capital expenditure for hardware.
- Solution: Leverage cloud platform environment services to improve availability, e.g., live migration from existing platform to target platform in case of outage.
- Benefits: As component re-hosting in multiple cloud platforms and improve availability and avoid vendor lock-in.
- Risks: Cloud providers do not provide the necessary services for applications to run in cloud platforms without re-architecting or rewriting code.

Migration paths emerge as sequential compositions of these patterns on a source architecture, see Fig. 1. These paths are defined based on discussions with the company about their existing architecture and a high-level specification of technical and business targets. Migration paths define decision points where typically several architectural options emerge, e.g., different data storage options (Xiong et al., 2013; Pahl and Xiong, 2013; Pahl et al., 2013). For (a subset of) these options, an experimental evaluation can be considered. Based on these paths, a plan focusing on a subset of components is identified for experimental evaluation: Firstly, define source and possible target architectures. Secondly, select critical components, e.g., high volume data process to test scalability of storage (DB) or communications infrastructure to test integration and communications scalability. The benefit of the patterns is that they link architecture configuration to quality. We use this link to select components for the feasibility exploration based on the most relevant quality concern to be explored.

3 EXPERIMENTAL MIGRATION FEASIBILITY STUDIES

Experimentation is the central activity during the feasibility study. Its objectives are:

- QoS and Cost: Quantification of QoS and cost aspects. Often, scalability is an important concern, driven by the business objective to expand. In this case, an experimental feasibility study can validate a proposed architecture scalability. Another motivation is a cost-vs-performance experiment, i.e., to consider different options and compare them technically (e.g., scalability of different target architecture options), but rank them considering the costs they would entail (Pahl et al., 2017; Fowley et al., 2013).
- Usage and Cost: Considering 'usage' changes the focus to the potential user. Usage exploration through experiments is a suitable tool to explore usage patterns and predict potential income based on this. This can then directly be related to the resources (and their costs) to facilitate user requests.
- Understanding: Experimentation allows to achieve a better understanding of technical constraints and operational activities in the cloud. What experimentation can shows is the difference between PaaS/IaaS/SaaS solutions (as consumer and provider) and integration and interoperation problems. It also clarifies how to structure and cost a staged migration (plan derivation).

An important question arises that links quality, usage and costs to the architectural configuration:

"How many processes can I host on a fixed cloud compute resource with a pre-defined latency performance target for a forecasted number of users of a particular application with a forecasted mix of application operation usage."

Experimental studies therefore play a central role in the migration feasibility determination. Experimentation often results in a prototype evaluation of a partly cloud-native architecture. Rather than just cloudifying a system in a virtual machine, we often selected a component such as data storage and have experimented with different cloud-native storage options, including for instance a mix of traditional RDBM and other table/blob storage formats. Partial experimentation with cloud-native prototypes allows

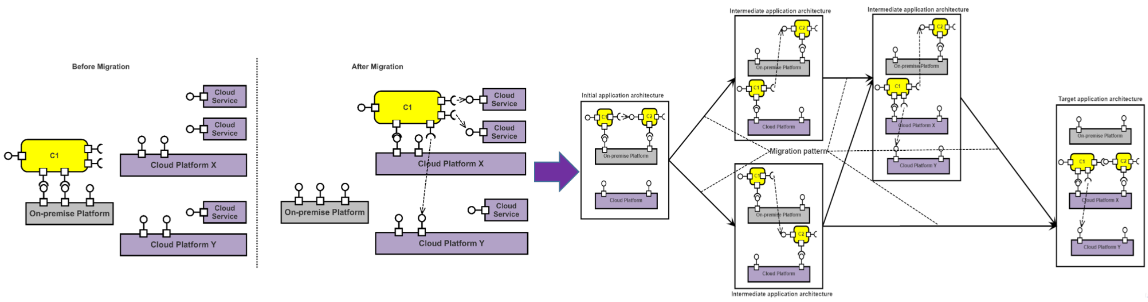

Figure 1: Cloud Migration Pattern (left) and a Migration Path based on several pattern applications (right).

to consider a fully cloud-native architecture to be discussed with realistic technical (e.g., scalability) and cost assumptions (storage, access) (Al-Room et al., 2013). Only realistic costs for cloud operation allows a charging model for their own product to be developed and validated.

When looking at concrete use cases, we distinguish what is expected and understood on the one hand and misconceptions and lack of knowledge on the other.

- Clarity of vision: Business reasons to go to cloud are often clear, e.g., a planned internationalisation or expecting an increase in company value (in the cloud). Technical reason to go to cloud are at a high level clear, for instance scalability.
- Understanding of cloud concerns (having an impact on architecture and process selection) is often limited. Technically, the difference between provision models/layers is unclear. This includes the management effort at I/P/SaaS level in comparison. Another problem is a possible vendor lock-in. In business terms, e.g., possibly required revenue model changes remain very unclear. Legal/Governance concerns such as data location are known, but without reliable knowledge.

We have carried out case studies in the following five application domains. We highlight the specific need for a feasibility study in each case:

- *Document management:* a document image processing to allow more efficient processing in the cloud, but also to enable the company to extend into new markets,
- *Banking solution:* an integrated solution (account management plus ATM operations) – provided in Africa and Asia, raising uncertainty concerns from security to legal,
- *Insurance:* a solution for multi-product management in multiple countries – uncertainty arise from the need for variability management of a single product across different regions/jurisdiction,
- *Food sector ERP:* an ERP solution for food production and sales – where a stable in-house solution is prepared for launch as a product into different markets. Food safety regulations impact on the architecture a cloud-based solution,
- *Business Registry:* enterprise repositories – scalable internationalisation is the driver, allowing clients to access their services through the cloud.

4 STRUCTURED FEASIBILITY EXPERIMENTATION

Experimentation aims to allow to establish a link between technical feasibility and quality (e.g., performance) and costs (Xiong et al., 2013). The pattern-approach helps to guide the select the components for exploration based on the most relevant quality concern. This can range from performance to security to integration and interoperability, as indicated for the use cases.

Furthermore, in many development approaches, quality testing is an integral component, in particular if the software is directly used by end-users. Load testing is difficult if prior testing has been done in-house and no expertise in testing in the cloud exists. The feasibility study thus also addresses load testing.

4.1 Process

First, we need a process for the feasibility study:

1. define source and possible target architectures,
2. select critical components, e.g., a high volume data process to test scalability of storage (DB) or a communications infrastructure to test integration or communications scalability.

Challenges for the migration expert and the company architects are to understand the existing architecture and to select best component for experimentation (ranges from individual component to full virtualisation as VM). The feasibility process is based on architecture determination, prototype component selection, and construction of realistic use case: (i) setup of services and data (real or dummy) and monitoring, (ii) experiment specification, (iii) experiment execution, (iv) data collection and analysis.

4.2 Experimentation Benefits

Sample experiments that we carried out targeted e.g. data storage for image processing in the 'Document Management' use case. What experimentation shows are a number of concerns that would normally not always be identified and clarified in discussions and non-experimental analyses:

- Difference between PaaS/IaaS/SaaS solutions (as consumer and provider). Based on the migration architecture it allows to practically demonstrate configuration and operation of different scenarios.
- Scalability of the solutions, as exemplied in Figures 2 and 3 for response time behaviour based on different retrieval and update loads.
- Identification of integration and interoperation problems that emerge in the prototype implementation when on-premise components are migrated.
- How to structure and cost a staged migration (plan derivation) can be estimated on the prototype component migration that typical involves the migration of existing components. In order to experiment, data and traffic needs to be available, again either imported or generated. In all cases, a realistic prediction of software and data migration tasks and related costs is possible.

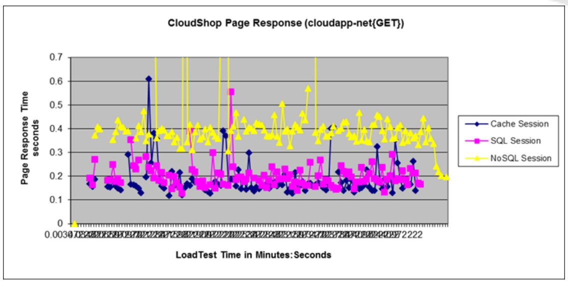

Figure 2: Comparison of storage options (technical quality).

Fig. 2 shows the performance results for different storage options for a Web application. Fig. 3 maps the architectural options onto costs.

4.3 Justification of Experimental Effort

Cost is a key factor to decide whether to conduct a feasibility study, i.e., does it pay off to carry out a feasibility study. The costs include the cloud platform to experiment with and the migration expert.

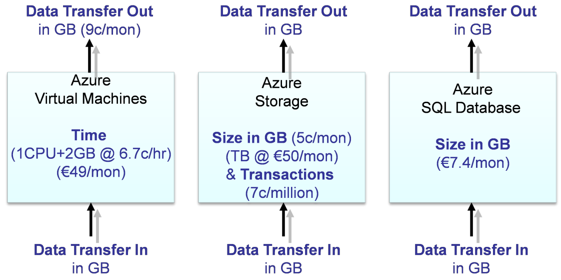

Figure 3: Comparison of storage options (data consumption-versus-cost).

The total experimental costs are 5-10% of the predicted migration costs based on our case study experience – including architectural analysis, draft migration plan and security and data protection analysis[1].

5 OBSERVATION & EVALUATION

The feasibility benefits are experimentation results and documentation: (i) a proper documentation of scenarios and plans needed at a high level, which allow detailed migration plans for a full migration to be developed, and (ii) experiments help to clarify options, address misconceptions, and identify open problems. They help to scope a full migration project. We have evaluated the feasibility approach to cloud migration based on the following criteria:

- Clarification of architectural options ans concerns
- Reliable quality prediction & cost prediction
- Cost effectiveness as an instrument

5.1 Clarify Architectural Concerns

Clarification of architectural options for SME through architectural prototyping using a pattern approach allows to use patterns that link quality to structure (Pahl, 2005; Pahl and Lee, 2015). A survey has been carried out with architects involved in the migration studies – with at least 3 architects for each participating use case). This includes architects both within the company in question and also the architects from the consulting organisation.

The important survey results are as follows: All participants (100%) agree or strongly agree that the method is suitable as an analysis tool to identify options and concerns. 88% agree that the migration method is suitable to analyse and discuss functional and non-functional architecture requirements for migration. One limitation has been flagged. Whereas

[1] These cost were part-funded for the given use cases by governmental business innovation schemes.

55% strongly agree that the method is suitable for SMEs and that is also suitable for multi-cloud migration (80% positive), 43% have concerns with its applicability for large-scale migration.

5.2 Predict Quality and Cost

We distinguish quality and cost observations:

- Reliable Quality Prediction: experimental quality assessment results in trustworthy, reliable input to predict full system behaviour. This has been discusses and analysed with the customer company. In the case of successful migrations at a later stage, we have not found any major deviations from the predictions when using the configurations experimented with.
- Reliable Cost Prediction: experimental cost assessment provides a low-to-high demand cost range, allowing to predict to system costs on a reliable basis (Wang et al., 2012). These have been discussed with the customer company and, where possible, evaluated through a later full implementation.

Only an empirical evaluation of the two quality items is currently possible due to the lack of suitable quality and cost mapping models. Cloud computing is here, however, highly suitable as monitoring of technical details can be set to be detailed for the experiments and billing for the resources used is equally detailed with typically metered usage for different resources used. For instance, for the 'Document Management' use case, the factors considered for storage only were: *region, replication options, data size stored, transaction number, data transfer*.

5.3 Cost Effectiveness

Cost effectiveness as an instrument is a key aim of our approach. We look at the actual costs of feasibility studies relative to full migration costs. We also consider how successful feasibility studies were in terms of determining the actual migration.

- Cost per project: Across the 5 migration cases that we have carried out with SMEs, the average cost was around 5% of the total budget for a full migration. The full migration budget was determined after the feasibility study concluded. The 5% cost was considered adequate. Please note that the overall budget for the migrations ranged between 100,000 and 250,000 Euros, including substantial re-architecting in some cases.
- Decisions taken: Concrete results from the use cases are as follows: 1 full migration (document management) has started and is currently being finalised, 1 full migration (insurance) has started, 1 decision against due to quality grounds (banking), 2 decision have at this stage not yet been taken (food, registry). In the all finalised cases, the companies in questions have based their decision on the outcome of the feasibility studies.

For the discussion with the companies in question, readiness to make a decision to embark on a full migration is considered a success. The experimental feasibility study is also considered successful if a decision against migration is taken on the grounds of an unfavourable quality or cost result.

5.4 Discussion

Limitations of the approach are related to the delay it might cause. In many circumstances this might not be an issue, but in some situations an approval to embark on a full migration is necessary, possibly depending on a successful outcome of the study. This approval is typically one of the following:

- External private investment in a cloud
- Public support, e.g. through innovation schemes
- Internal approval for further funding

While there might be a delay in these situations, any decision can be based on reliable input data. Note that in the use cases where a decision has not been taken yet, this was not due to the study results.

6 RELATED WORK

A number of migration approaches exist. Jamshidi et al. (Jamshidi et al., 2013) survey the literature. Fahmideh et al. (Gholami et al., 2016) provide a comparison framework for cloud migration. We target specifically an experimental framework that goes beyond process models for migration.

Costing models for cloud are important for SMSs to understand their own costs and expenses. In (Arshad et al., 2015), an overview of pricing models for the cloud is provided. On the expenses side, e.g., at IaaS level, resources are priced often like commodities (Wang et al., 2012). At the income side for the SME, the product is typically provided as a SaaS with possibly a mix of model from pay-per-use to pay-per-user and flat-fee models. Our solution is meant to bridge between the two perspectives.

Quality in cloud is managed for factors as performance by configuring the resources used appropriately. We propose a manual experimentation approach for cloud prototype implementation. In (Son,

2013) a system to support automated resource selection is suggested. Although this is not generally application, the ideas could help to automate our approach further. The automation of cloud experimentation is also addressed in (Affetti et al., 2015) through a tool suite for OpenStack.

7 CONCLUSIONS

Reliability of data and an understanding of the processes and their impact are essential for any decisions (Chappell, 2016). For a cloud-based software provider, quality and cost need to be reconciled in an architecture that maps IaaS or PaaS hosted software onto a SaaS delivery model.

We suggest feasibility studies that clarify the costs of both migration and also the operation of a software product in the cloud. They provide decision support, allowing to answer whether/how a software product can be deployed and delivered cost-effectively in the cloud while maintaining required quality. The benefit is increased reliability of data/assumptions, rather than relying on experience or guesses.

Experiments have another benefit. They cover tasks normally not done until a systems deployments stage. Performance engineering is important, but load tests are normally not done until a late project stage. Here, load testing is a key experimental focus that allows to reduce technical risks at a very early stage.

As already mentioned, one aspect for future work is the increased automation of the experiments. This could include automated test case generation for scalability tests or even the selection of different alternative services for a given component (e.g., an automated storage service selection and configuration).

ACKNOWLEDGEMENTS

This work was partly supported by IC4 (the Irish Centre for Cloud Computing and Commerce), funded by EI and IDA.

REFERENCES

Jamshidi, P., Ahmad, A. and Pahl, C. (2013). Cloud migration research: a systematic review. IEEE Transactions on Cloud Computing.

Jamshidi, P., Pahl, C. and Mendonca, N.C. (2016). Pattern-based multi-cloud architecture migration. Software: Practice and Experience.

Son, J. (2013). Automated Decision System for Efficient Resource Selection and Allocation in Inter-Clouds. The University of Melbourne.

Arshad, S., Ullah, S., Khan, S.A., Awan, M.D. and Khayal, M. (2015). A survey of Cloud computing variable pricing models. In Eval of Novel Appr to Sw Eng.

Jamshidi, P., Pahl, C., Chinenyeze, S. and Liu, X. (2014). Cloud migration patterns: a multi-cloud service architecture perspective. WESOA.

Xiong, H., Fowley, F., Pahl, C. and Moran, N. (2013). Scalable architectures for platform-as-a-service clouds: performance and cost analysis. Europ Conference on Software Architecture.

Pahl, C. and Xiong, H. (2013). Migration to PaaS Clouds - Migration Process and Architectural Concerns. MESOCA Symposium.

Pahl, C., Xiong, H. and Walshe, R. (2013). A comparison of on-premise to cloud migration approaches. Europ Conf on Service-Oriented and Cloud Computing.

Al-Roomi, M., Al-Ebrahim, A., Buqrais, S. and Ahmad, I. (2013). Cloud Computing Pricing Models: A Survey. Intl Jrnll of Grid and Distr Comp. Vol.6, No.5.

Wang, W., Zhang, P., Lan, L. and Aggarwal, V. (2012). Datacenter net profit optimization with deadline dependent pricing. Conf on Inf Sciences and Systems.

Giardino, C., Bajwa, S.S., Wang, X. and Abrahamsson, P. (2015). Key Challenges in Early-Stage Software Startups. XP Conference.

Li, H., Zhong, L., Liu, J., Li, B. and Xu, K. (2011). Cost-effective partial migration of VoD services to content clouds. Cloud Computing.

Fowley, F., Pahl, C. and Zhang, L. (2013). A comparison framework and review of service brokerage solutions for cloud architectures. ICSOC Workshops.

Pahl, C., Jamshidi, P. and Weyns, D. (2017). Cloud architecture continuity: Change models and change rules for sustainable cloud software architectures. Journal of Software: Evolution and Process.

Pahl, C. (2005). Layered ontological modelling for web service-oriented model-driven architecture ECMDA-FA, LNCS 3748, pp. 88-102.

Chappell, D. (2016). Cloud Computing White Papers. http://www.davidchappell.com/writing/white_papers.php.

Gholami, M.F., Daneshgar, F. and Rabhi, F. (2016). Cloud Migration Methodologies: Preliminary Findings. CloudWays Workshop.

Pahl, C. and Lee, B. (2015). Containers and clusters for edge cloud architectures - a technology review. Intl Conf on Future Internet of Things and Cloud.

Affetti, L., Bresciani, G. and Guinea, S. (2015). aDock: A Cloud Infrastructure Experimentation Environment Based on Open Stack and Docker. Intl Conf Cloud Comp.