Decision Support for Partially Moving Applications to the Cloud – The Example of Business Intelligence

Adrián Juan-Verdejo[1,2], Dr. Henning Baars[2]
[1]CAS Software A.G, Karlsruhe, Germany
[2]Chair of Information Systems I, University of Stuttgart, Stuttgart, Germany
adrian.juan@cas.de, baars@wi.uni-stuttgart.de

ABSTRACT

Cloud computing services have evolved to a sourcing option that promises a wide range of benefits, such as increased scalability and flexibility at reduced costs. However, many enterprise applications are subject to strict requirements – e.g. regarding privacy, security and availability – and are embedded into complex enterprise IT architectures with a multitude of interdependencies. For these reasons, many decision makers have developed a sceptical stance towards cloud computing. A solution might be a *hybrid* (local/cloud infrastructure) approach where only suited components are migrated to a cloud infrastructure. This, however, has significant architectural consequences that need to be taken into account. This contribution suggests a cloud migration framework that will be implemented as an IT-based decision support system based on modelling the interdependencies between components. The approach is illustrated with the example of Business Intelligence (BI), i.e. integrated approaches to management support. The underlying decision model would particularly consider data transfer volumes, performance, sensitivity of cloud-based data repositories, as well as exposure to public networks. The potential of such an approach is illustrated with a selected set of BI scenarios. Based on this, conclusions are derived and generalised for approaches taking into account deployments on both the local premises and cloud infrastructures.

Categories and Subject Descriptors

D.0 [**Software**]: General; H.3.4 [**Information Storage and Retrieval**] Systems and Software: Distributed systems;
J.1 [**Computer Applications**] Administrative Data Processing: Business

General Terms

Design, Security, Experimentation, Management, and Performance.

Keywords

Cloud computing, decision-making, multiple criteria decision making, migration, business intelligence, enterprise applications, and security policies.

1. INTRODUCTION

Enterprises are interested in using public cloud offerings mainly because of the potential to enhance their enterprise applications with increased scalability and flexibility at reduced costs [1]. Besides the financial and technical benefits stemming from the use of computing, storage and networking resources in public cloud offerings, some improvements arise from the organizational point of view, mainly regarding an increased agility and an enhanced division of labour [1-3]. The ability to develop, test, and scale applications in environments with volatile and unpredictable demand is a very interesting feature not only for rapidly evolving start-up companies [2] but also for companies pursuing innovative projects or seeking to react to sudden changes in the business environment. However, these benefits can easily be counterbalanced when security and privacy issues come into play and when interdependencies with other pre-existing applications gain in relevance and strength [4]. Additionally, despite the potential immediate cost reductions that can be achieved by running a single application or component in a cloud environment [5] (e.g. a JBOSS application server component or a web component for data visualization), the cloud-based provision might have negative effects on the overall lifetime cost of the higher-order enterprise system. For example, a component for data reporting could be migrated to a cloud infrastructure, and seemingly incur less costs. However, if we consider the BI system as a whole, the side effects on the overall system could cause increased costs because the data is not integrated or the reports are not consistent.

More specifically, there are technical, economic, and compliance-related factors that have to be considered when moving an enterprise application to a cloud environment. From a *technical* standpoint, performance requirements of an application, availability, traceability of errors, data and system lock-in problems, as well as interoperability issues have to be considered [2]. Applications to be moved to a cloud infrastructure often consist of a high number of diverse application and middleware components which depend on and interact with each other [4]. Due to this complexity, the migration decision regarding what tiers or even what components should be moved to what specific cloud infrastructure becomes difficult [4]. *Economically*, financial risks related to the cost of wide-area communications, the migration of large amount of data to cloud environments, or the correct estimation of resources pose challenges [6]. With respect to *compliance,* organisational, governance-related, contractual, and legal constraints have to be respected, particularly regarding security levels and data sensitivity. Between many of those factors there are some trade-offs – gaining one quality might imply losing on another aspect.

An option to cope with these problems – and the security-related problems in particular – is to move to a model where only non-sensitive data and computation is moved to cloud environments, whereas critical data remains locally at the client organization. In such situations, ideally, components dependent on those data might be forced to stay next to the data they depend

HotTopiCS'13, April 20–21, 2013, Prague, Czech Republic.

upon in order to avoid the introduction of wider-area communications or decrease responsiveness.

These issues are particularly visible in the field of integrated management support – Business Intelligence (BI). BI usually involves complex multi-layered architectures, contains highly sensitive and strategically relevant contents, is subject to agility requirements, involves large volumes of data to be stored, transferred, and processed and is subject to fluctuating workloads [7-10]. Due to the complexity associated with moving BI systems to cloud-enabled environments, we argue for the need of assisting the decision maker during this migration process. Moreover, we strive for generalisation so that our approach is applicable to more application settings rather than just BI systems.

To address the above mentioned challenges, we propose a cloud migration framework taking a decision support approach for the construction of architectures combining the local premises and cloud infrastructures and we come up with a first draft that is based on an iterative and systematic selection of cloud-based components. The approach is illustrated within the field of Business Intelligence.

2. RELATED WORK

In this section we describe related work on cloud computing that has been put into context in moving applications to cloud environments. We present different research works related to the distribution of components on both cloud and local premises. We also include in the discussion approaches related to the selection of the best cloud services provider for a specific combination of parameters as the problems are similar in structure (choice of system distribution options). The division and scattering of the target application is driven by many technical factors with various business, organizational, economic, security and privacy implications [1, 4]. With the aim of facilitating the comparison of the presented research works and their different focus, granularity and application settings, **Error! Reference source not found.** compares the selected research works [4, 6, 11-17] in terms of the specific focus or point of view taken for the migration of applications to cloud infrastructures and the applications' settings the approach applies to (e.g. MapReduce applications, Web applications or more general approaches). In the third column, **Error! Reference source not found.** shows at what level of granularity the migration is taking place. The factors to be taken into account before performing the partition of an application at the level of code entities are quite different from the case where entire components are migrated.

The partitioning of applications has been extensively explored for client-server architectures [18, 19]. However, the settings for application partition in cloud environments differ from the server and client partitioning case. Particularly, the effective design and partitioning of legacy enterprise applications is not as obvious [13]. Several research works [6, 12-14, 17] have been done to assist the decision making process in order to withdraw this responsibility from applications' suppliers — either by taking a general approach for the development of frameworks valid for a wide set of applications or a finer grained solution for specific infrastructures or applications' settings. Such a framework for alleviating the migration burden from the user is Manticore [12].

Research work	Focus	Granularity	Application's Setting
Moving-to-the-cloud problem[4]	Moving applications to cloud environments	Components	Pre-existing applications
Cloudward Bound [6]	MCDM migration	Components	Enterprise Applications
Volley [11]	Data partitioning	MapReduce Jobs	MapReduce
Manticore [12]	Code partitioning	Code entities	Software services
HybrEx [13]	Data and system partitioning focusing on privacy	Components	Distributed Applications
COPE [14]	Automated orchestration using declarative languages	Virtual Machines	Distributed applications
CloudGenious [15]	Web servers MCDM* migration	Virtual Machines	Web Applications
$(MC^2)^2$ [16]	MCDM migration	Virtual Machines	Generic (conceptual framework)
Conductor [17]	Orchestration deployment	MapReduce Jobs	MapReduce

Table 1 Comparison of the different relevant research works
MCDM = Multiple criteria decision making*

Kaviani et al.'s Manticore is a framework for the partitioning of software services for a deployment on the cloud and local premises. Manticore tries to automate to some degree the application partitioning by using two application dependency models the authors found in literature, namely the request-based model and the static structure model. The former decides whether to execute a request fully either on the private or the cloud premises, whereas the latter provides the possibility to partition a request between both of them. Additionally, Kaviani et al. present the context-sensitive model in order to address the problems they found in the previously mentioned models. By using this model a request can be partitioned to be executed both on a cloud environment and the private premises while keeping a differentiated transitive set of callers to each function execution (Kaviani et al. call it *calling context*). Their research coincides with our approach as we both fix our attention in the software services partitioning. Nevertheless, their level of abstraction is lower than ours as we consider moving components rather than code entities, i.e. functions. This is more realistic for enterprise systems that are often built with commercial packed software components. An example is the domain of Business Intelligence where systems are developed by customizing and integrating commercial tools for (among others) reporting and analysis, DWHs and data marts, as well as for ETL and data management components (from large vendors like IBM, SAP, SAS, Microstrategy, Oracle, Teradata, or Informatica). While Manticore focuses only on the partitioning of software services, some research efforts, such as Volley [11] focus on data partitioning. They aim at increasing performance and reducing traffic at the datacenter by performing data relocation rather than partitioning the code like in Manticore [12]. Although our approach considers the data partitioning, as a first step it is based on an all-or-nothing data partitioning strategy.

Some other research works suggest partial migration of specific types of applications to cloud infrastructures at different levels of granularity. Some research works focus on the migration of web applications. Menzel and Ranjan's CloudGenius framework [15] supports the decision-making involved in the migration of Web servers to cloud environments. CloudGenius takes into account how a large set of heterogeneous criteria and their interdependencies affect the migration of multi-component enterprise applications. Nevertheless, they apply their framework to single-tiered applications and leave out a lot of current enterprise applications which are more complex than that. CloudGenius selects the best and compatible mix of software images (e.g. Web server image) and infrastructure services to ensure the targeted Quality of Service (QoS) of the application. Menzel and Ranjan leverage the Analytic Hierarchy Process (AHP), rather than the Analytic Network Process (ANP) suggested by their previous work $(MC^2)^2$ [16], to enhance the multiple criteria decision making (MCDM). Our approach also leverages AHP to assist the migration decision.

Similarly to CloudGenius and $(MC^2)^2$, some other research efforts try to enhance the selection among the available cloud offerings according to cloud computing customers' requirements. Hajjat et al.'s Cloudward Bound [6] focuses on the assessment of the beneficial migration of enterprise applications to cloud infrastructures following a deployment both on the local and the cloud premises. Hajjat et al. show with Cloudward Bound how the overall migration cost is the result of a complex combined effect of applications characteristics in terms of workload intensity, storage capacity, growth rate, and the cost of software licenses. They conclude that horizontal partitioning between in-house and cloud deployments can be very beneficial for certain applications. However similar to our approach, they do not consider the cloud offering selection in their approach. We consider this an important feature as the cloud offering selection determines the applicable cost model and affects the provided QoS and non-functional requirements. Cloudward Bound partitions and relocates components at the level of application servers or virtual machines. Likewise, Liu et al.'s COPE [14] aims at the partitioning (local and cloud premises) of computation and data at a high level of abstraction. With COPE, Liu et al. orchestrate a set of subsystems (computing, storage, and network resources) that span large geographic areas serving different clients with the aim of cost optimization on a distributed setting [14]. Our approach differs from these research works because we consider a holistic migration strategy incorporating the generation of different migration alternatives and the selection of the more appropriate one.

Other research works are concerned with another sort of applications, namely MapReduce applications. Such is the case of the work done for Conductor [17] and Hybrex [13]. As for HybrEx [13], the authors propose a new execution model for confidentiality and privacy in cloud computing. With HybrEx, Ko et al. outline data, system partitioning, and the integration of both. We adopt some of the ideas in HybrEx but at the same time consider more factors affecting the migration in addition to security and confidentiality. Further, Conductor [17] integrates an abstraction layer into Hadoop [20] in order to help cloud computing customers select the right set of resources to meet their requirements. Furthermore, their design includes mechanisms to use different services seamlessly and frees the cloud computing customer from manual choice and optimization. These two research works relate to ours but we aim at enabling the migration of a larger set of systems as we do not only pay attention to MapReduce applications.

Our research work's novelty lies in the holistic approach taken for the partial migration of applications to cloud infrastructures by taking into account many interdependent parameters. These parameters include business and economic considerations, technical and security-related challenges, and the organizational implications of the partial adoption of cloud computing as computation model. Further, our work incorporates the selection of the right cloud provider given the application's requirements.

3. MOVING SYSTEM TO CLOUD ENVIRONMENTS: *The example of cloud-based BI*

With the aim of evaluating our approach for partially moving enterprise systems to cloud infrastructures, we focus on migrating BI systems to a cloud environment. *Business Intelligence* denotes integrated approaches to IT-based management and decision support. Current BI landscapes are usually built upon multi-layered architectures. The foundation is the Data Warehouse (DWH), an integrated, application-spanning data repository that contains data extracted from a variety of internal and external data sources and usually feeds application specific data pools, *data marts*. These are accessed by a variety of analysis and reporting applications ranging from OLAP and data mining / predictive analysis tools, to planning, budgeting and balanced scorecard applications. The components are bound together by administration and metadata tools [21-23]. Under the heading *operational BI*, the systems have also been extended to the support of operational decisions, e.g. in the realm of process management [24, 25]. The complementary analysis of large volumes of structured and unstructured data (especially from the Internet), which is discussed under the term *Big Data* [26, 27], also becomes increasingly relevant for BI. In many companies, specifically designed BI governance frameworks have been implemented in order to govern the resulting architectures, processes, and applications [28].

While some BI applications do not contain much sensitive information (e.g. for web usage mining), others deal with strategically relevant data, or need to be aligned with various legal requirements, e.g. in the fields of (legal) reporting or the analysis of user data.

From this perspective it is clear that a complete BI system cannot easily and completely be *outsourced* to cloud environments. On the other hand, cloud offerings might ease existing agility pressures [9] and be a response to both volatile capacity requirements (e.g. in planning and budgeting) as well as for project-specific BI solutions (e.g. for social media analysis) [29-31]. Both market studies [32, 33] and the wide spectrum of BI SaaS services, from almost all large BI vendors, indicate that there are indeed high expectations towards bringing BI to cloud infrastructures. Consequently, BI systems are not only a relevant area for a cloud-based application provision but also one where a deployment model distributed both on the local and cloud premises is of particular relevance. The following scenarios illustrate the benefits of our approach. They are derived from [29] as well as from other research scenarios we have investigated.

Scenario 1: Inclusion of specialized data analysis functionality from cloud infrastructures, e.g. for Big Data analysis

In this scenario, specialized analytic functionality is sourced from cloud infrastructures, especially for temporary projects, e.g. when conducting a market analysis based on web site and social web data with Big Data approaches. In this case, the envisioned software would primarily consider the interdependencies with existing *data repositories* (data marts, DWH) and the underlying DBMS. As long as the analysed data is non-sensitive, the envisioned decision support software would primarily rank the options of where to store the data based on the initial volume of the data that has to be transferred to a cloud environment with the result probably being the movement of the data mart to a cloud-based environment - which is selected primarily based on cost rationales.

Scenario 2: Movement of a reporting or OLAP frontend to cloud computing environments

The second scenario is less straightforward and illustrates the benefits of a tool-based approach: Moving a reporting and/or an OLAP frontend to a cloud environment — i.e. separated from the rest of the BI landscape — can have clear benefits, e.g. regarding accessibility with mobile devices. However, it also has implications for reaction times to the user's input and data traffic. Dependencies e.g. exist with the Device and User Management, the Web and the Application servers and most importantly: the relevant data repositories. The movement of the data mart (and the DBMS) in particular might alleviate the reaction time issues but has further implications regarding security (not only a few reports are exposed to the web but the whole data repository) and consistency with the DWH. In the case of operational reporting, the continuous updates of the data mart in a cloud environment might induce prohibitive bottlenecks in cost and performance – leading to a solution where the DWH (or a specifically designed *Operational Data Store*) also needs to be brought to a cloud infrastructure.

Scenario 3: Movement of an Operational BI solution to the web

The dependencies get even more critical when BI software triggers events in other systems (Active BI), e.g. on operational level: here, it needs to be decided where the functionality that triggers the event is kept. Also, like in Scenario 2, operational data needs to be updated in real time – but in this case the dependencies go in both directions. The dependencies here also include transaction management and operational source (or *destination* systems).

Scenario 4: Movement of selected ETL procedures to cloud environments, particularly for data sources that are already in the cloud infrastructures

There are reasons why including specialized ETL routines from a cloud environment might be an interesting option, e.g. for pre-processing unstructured data or for discovering non-evident duplicate entries (e.g. sales for Müller and Mueller which in the end are the same company) in the master data that is fed into the DWH. In this case, the routines have to be embedded into the higher-order ETL process. Given the fact that this often involves really large amounts of traffic and that the ETL process is a link between several core components, this will affect the overall BI system on various levels.

4. APPROACH AND SCENARIOS

The decision-making process related to the migration of an existing application to a virtualized cloud system is assisted through the use of our proposed cloud migration framework. In this section we describe our approach in providing the framework. The case of the migration of Business Intelligence systems to cloud infrastructures is shown as an example of a system to be migrated. Our approach tries to fill the gap we found in literature by considering the multiple factors that affect the migration to virtualized cloud environments instead of focusing in a subgroup of them. We call for a cloud migration framework which takes a holistic approach for migration that also considers the particularities of the cloud offering selected to run the application and the interactions with the local data centre.

Firstly, our approach lets the decision maker describe the previously existing application (or system) in terms of its components, the dependencies between them, and the specific characteristics of those components. Next, several migration strategies are generated with a model for migration to cloud environments in mind where some of the applications' components are kept locally while others are migrated to cloud-based infrastructures. Further, a migration strategy is selected based on its compliance to the requirements selected and ranked in importance by considering the characteristics of the suitable cloud offerings where the application could be deployed and the input from the decision maker captured within the Analytical Hierarchy Process.

4.1 Architectural description of the pre-existing system

As a first step, we envision that the pre-existing system targeted for migration to a cloud infrastructure is described in terms of its architecture (*see* example in *Figure 1*). By using our framework the system's architecture is described and extra information is added to the model. The results obtained after this step serves as input for our proposed cloud migration framework. Our framework generates several migration alternatives given the architectural description of the pre-existing application. The system's architectural description is annotated with some additional information. These annotations will serve as a basis for the estimation of how the system will perform after migration and will therefore affect the design of the migration alternatives.

The details of the pre-existing system are specified including both the application's *building blocks* and *users*. The *building blocks* are defined as components and the tier they belong to (e.g. *back-end*, *business-logic*, *front-end*). The components can either be application components, such as UIS, services, workflows, databases, or middleware components, such as application servers, workflow engines, or database management systems (DBMS). Furthermore, the *users* are described regarding whether their location is internal or external to the deployment. The relation between different components as well as between components and the application's users are modelled with the aim of explicitly stating the application's dependencies (*see Figure 1*). In the system's architectural description (**Figure 1**), the Rack, JBOSS 4.2 and JBOSS 6.0 Application Servers – and the components they consist of – are shown. As for the red arrows they show the interactions happening within the system. The system in **Figure 1** access data stored in two components, namely the *operational data store* and the *management info* components. As it is shown in **Figure 1**, the user only interacts directly with the *BI-Info-Presenter* component on the right part of the figure.

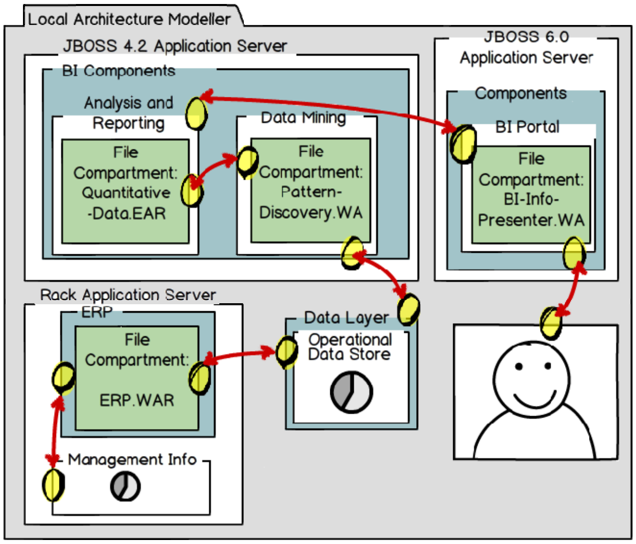

Figure 1 Pre-existing System's Architectural Description

In addition to the architectural description of the system, our migration alternatives generator needs input regarding the nature of the interdependencies among components and how the transactions between them happen. With this information at hand our framework can better estimate the effects of the migration and recommend a migration alternative that will better fit the system.

As an example, we show in *Figure 2* an example of annotations to the architectural description presented in *Figure 1*. In the example shown in *Figure 2* a transaction – between a component for data mining and a persistence component within the data layer – is annotated in terms of the origin and destination of the relation (at the top of **Figure 2**) together with the information related to the frequency of the transaction, the amount of data transferred, and values related to the delay, e.g. average delay, its variance, or the peak delay value registered (at the bottom of **Figure 2**).

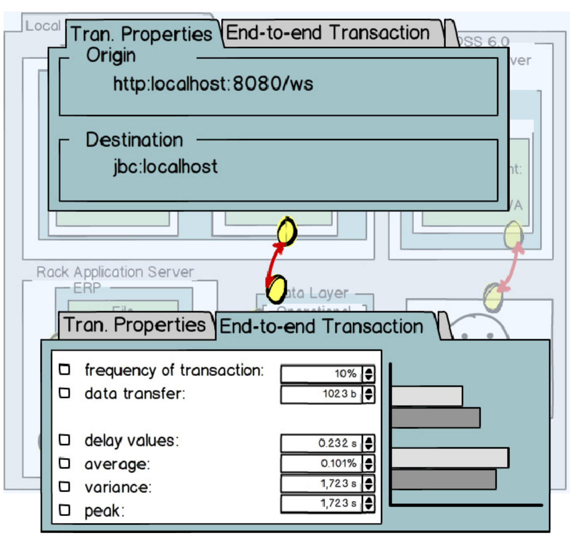

Figure 2 Pre-existing System's Architectural Description Annotated

For our approach, it is important to describe the transactions between components but it is also crucial to consider legal and regulatory constraints as well as functional and non-functional requirements. This information will enable our proposed framework to be more accurate in the estimation of how the target system will perform after migration while complying with the users' requirements.

There are several requirements which will have a big impact on the migration alternative generation. An example of a functional requirement that might affect the migration process is whether a component needs to be run in a specific middleware (e.g. *JBOSS*). As a result, the target cloud offering for that component is required to be able to run that middleware. As for the non-functional requirements, the proposed framework will consider information related to data privacy (e.g. *what* can be *where*), cost, or offered Quality of Service (QoS).

4.2 Migration strategies generation

Our cloud migration framework takes into account the input from the previous phase to generate several complying migration alternatives. The migration alternatives follow a deployment on both the cloud and the local premises in order to cope with security and legacy issues associated with the migration. This deployment benefits from the possibilities offered by cloud computing while respecting the fulfilment of the SLA and the respect to policy constraints.

The migration alternatives are generated based on a model for component placement. The goal is to bring the application to a cloud computing environment with the aim of maximizing the monetary benefits as well as respecting the functional and non-functional requirements defined in the previous phase. Let us consider a system S with several internal and external (to the local data centre) users U_w, $1 \leq w \leq m$. The users are clustered in these two groups as an internal user would be preferably routed to the local data centre, whereas an external one might better fit to an interaction with a component in a cloud infrastructure. The system described within the previous phase (*see* Section 4.1 and **Figure 1**) runs several applications, each composed of several components, C_i, $1 \leq i \leq n$ that communicate with each other. The transactions within the system are described in a transactional matrix $TM = \{C_k, X_m\}$, where X_m is either a user or another component. The transactional matrix is annotated with the frequency and data exchange per transaction.

The output of this phase will consist of several migration alternatives, M_r, $1 \leq r \leq p$, where M is a set of components C to be migrated to a cloud environment. In the end we face a maximization problem:

Max	Benefits (M) – InternetCosts(M)
Subject to	Policies Fulfilment of the performance requirements Respect to pre-existing interdependencies

Benefits(M)

The benefits of moving a system to cloud-based infrastructures is expressed as the difference between how much it used to cost to run a system and how much it costs after the migration. The benefits stem from the fact that large data centres are able to offer cheaper computation, storage and networking. These large data centres benefit from being economies of scale and let the user request resources on-demand. Therefore, the user does not have to incur higher costs by over-provisioning. Furthermore, moving applications to a cloud environment lowers the operational expenses due to the outsourcing of the IT department. Applications prone to benefit from cloud computing are those which involve either computing or storage-intensive tasks (e.g.

Business Intelligence applications; particularly in cases as described in *Scenario 1* within *section 2.2*). The estimation of the economic benefits due to migration is not easy and depends on the pricing model of the cloud environment where the application will be finally deployed. In any case it is one of the main drivers of the migration alternative generation.

InternetCosts(M)

As a result of the re-deployment, the system will incur increased communication costs due to the introduction of wider communication links between components which used to be at the same location. Due to the fact that those components leverage the Internet in order to communicate between remote locations, i.e. the local data centre and the cloud premises, the increased internet costs between them are factored in our model (particularly visible in Scenarios 3 and 4). These costs have to be included in the model as they will counterbalance the economic benefits due to the migration to a cloud infrastructure (*Benefits(M)*)

Policies

The resulting migration alternative has to respect the security requirements of the user and respect the privacy of sensitive data. In fact this is one of the reasons that lead us to advocate for a deployment on cloud-based infrastructures and the local premises in order to respect the data sensitivity, enterprise-specific constraints, industry-specific laws and regulations, and national privacy requirements.

Fulfilment of the performance requirements

Due to the changes in the placement of components, the overall behaviour of the target system may change but the QoS requirements of the user have to be met. Once some components are migrated to a cloud infrastructure, some delays might appear as those components interact with some others which remained locally (at the local data centre) in order to get some computation or data from them. Those delays appear because a wide-area link has been introduced between components which used to interact within the same location (as seen in Scenario 2). Let us consider that those delays trespass a given threshold, some components. In this scenario, most likely some components could not be moved to the cloud infrastructure in order not to damage the QoS.

Respect to pre-existing interdependencies

The logic of the system and the interdependencies have to be nevertheless maintained after migration. That is, if a component had some dependency with another one, this relation somehow has to exist in the new deployment after migration.

4.3 Deployment Options Ranking

The several migration strategies generated in the previous phase are ranked based on customer's requirements and the features offered by the different cloud offerings regarded as options for deployment. The different deployments on different cloud offerings present a myriad of Key Performance Indicators (KPI) with many attributes and sub attributes. As a result, deciding what option is the best-fitting one is a complex task due to its multiple criteria decision-making nature [34]. An additional difficulty originates in some attributes which are rather qualitative and therefore it is not immediate to give them a numerical value. Security and trust are typical examples of this kind of attributes.

The ranking mechanism we propose is based on a well-known process to solve these complex multiple criteria decision-making problems, namely the Analytic Hierarchy Process (AHP) [35]. We propose to leverage AHP in order to create a hierarchy structure of attributes that will lead to the identification of the best migration alternative according to the consumer's requirements. The first step within the AHP is to decompose the problem and select the relevant criteria. We model the problem in three levels (bottom-up): (1) alternatives to achieve the decision goal, (2) the several criteria applicable to evaluate those alternatives, and (3) the decision goal itself (*see 4.3.1)*. The second step is to establish priorities between the criteria of the hierarchy by using pair wise comparisons between those elements in order to help decision makers to judge what the best option is. As a result of this step the criteria will be ordered depending on the assigned priority (*see 4.3.2)*. Next, the priorities of the criteria in the hierarchy are combined with the different values given to those properties for every alternative (*see 4.3.3)*. Finally, the consistency of the judgments is checked and a final decision is taken (*see 4.3.4)*.

4.3.1 Problem decomposition and selection of criteria

Figure 3 presents a generic example of the AHP hierarchy. On the upper layer the main goal is stated. In our case it is to find the better strategy fitting the requirements described by the decision maker. On the bottom layer the different migration strategies generated by our framework (*see section 3.2*). The intermediate layer contains the different criteria and attributes related to moving an application to cloud environments. Some criteria are essential requirements which must be present in the elected alternative whereas others are rated depending on their importance for the system targeted for migration.

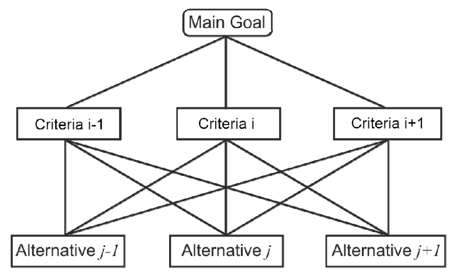

Figure 3 AHP hierarchy

4.3.2 Criteria prioritization

In order to be able to compare the different migration alternatives at the bottom, we need to assign weights (i.e. relative importance) to the criteria on the second layer. Pair wise comparisons of the criteria on the second layer are presented to the decision maker so that a criterion can be rated relative to another one. As suggested by [35], the rating is done in a 1 to 9 scale, from *equal importance* to *extremely more important*. Additionally, a user can arbitrary assign a weight to a criterion but afterwards the criteria will need to be normalized in order to let them be comparable.

4.3.3 Scoring

The migration alternatives in the bottom layer will have a numeric value for the criteria prioritized in the previous phase. Sometimes it is not possible to quantify some attributes and therefore the relative weights will be assigned to the migration alternatives based on the literature [36]. Firstly, the units of both values have to be the same to make them comparable. Further, different comparison metrics will be offered in order to compare those non-quantifiable attributes.

4.3.4 Decision making

From the previous phase, the scores for each criterion are combined with the relative importance of those criteria relative to the others and the cloud offering characteristics [37]. For every attribute on the second layer (*see Figure 3*), the combination of priority and value is aggregated for every alternative for migration. Finally, that alternative ranking higher will be picked.

5. CONCLUSION

Although cloud computing offers many benefits, the migration to a cloud-based architecture needs to be carefully planned. Decision makers interested in migrating their applications to cloud-enabled environments need to take many criteria into account in order to adopt an adequate migration strategy. This paper calls for a cloud migration framework to assist in the moving of the target application by following a local and cloud deployment model instead of an all-or-nothing approach. Our framework envisions that parts of the application are kept locally while others parts are migrated to cloud infrastructures. The decision making process of how to move the application to cloud environments is supported by a model for component placement and the implementation of the Analytical Hierarchy Process. Finally, both the whole framework and models are applied to the case of migrating BI applications to a cloud offering.

Due to the sensitive nature of the data within BI systems, we consider that a BI system cannot be easily and completely moved to a cloud environment. We envision four different scenarios in which BI applications are moved to cloud infrastructures. We present the scenarios and explicitly state what parts of the system are outsourced and what parts are not. The challenges and opportunities associated to the scenarios are shown. We explain how BI solutions can leverage cloud computing for the cost-efficient execution of computing and storage-intensive tasks.

This is an on-going research work which, as a result, will bring a refined version of the cloud migration framework that is implemented prototypically and – on these bases – evaluated in an industry setting. The next steps of our research are to clearly identify the criteria driving the migration decisions for every kind of system and for Business Intelligence systems in particular. Based on those criteria the model will be implemented and utilized to generate optimal migration alternatives. Further, we will explore the cloud selection process and incorporate cloud offering criteria. Finally, we will further structure and detail the BI scenarios in order to apply our cloud migration framework. This last step will serve as evaluation phase.

6. ACKNOWLEDGMENTS

This research work has been supported by *RELATE* Initial Training Network Marie Curie Action (ITN) within the 7th Framework Programme for Research and Technological Development (FP7) funded by the European Commission.

We would also like to thank Seyed Vahid Mohammadi and Dr. Bholanath Surajbali from CAS Software AG for their input, help and feedback.

7. REFERENCES

[1] Armbrust, M., et al., *A view of cloud computing,* Communications of the ACM, 2010, Vol. 53 (4), p. 50-58.

[2] Khajeh-Hosseini, A., et al., *Decision support tools for cloud migration in the enterprise*, in IEEE International Conference on Cloud Computing (CLOUD), 2011: IEEE.

[3] Rygielski, P. and S. Kounev, *Network Virtualization for QoS-Aware Resource Management in Cloud Data Centers: A Survey,* Praxis der Informationsverarbeitung und Kommunikation, 2013, Vol. 36 (1), p. 55–64.

[4] Leymann, F., et al., *Moving applications to the cloud: an approach based on application model enrichment,* International Journal of Cooperative Information Systems, 2011, Vol. 20 (03), p. 307-356.

[5] Rizou, S.a.P., A., *Towards value-based resource provisioning in the Cloud*, in IEEE 4th International Conference on Cloud Computing Technology and Science, 2012.

[6] Hajjat, M., et al., *Cloudward bound: planning for beneficial migration of enterprise applications to the cloud*, 2010: ACM.

[7] Moss, S.T. and S. Atre, *Business Intelligence Roadmap – The Complete Project Lifecycle for Decision Support Applications*, 2003, Boston: Addison-Wesley

[8] Turban, E., R. Sharda, and D. Delen, *Decision Support and Business Intelligence Systems*, 9th edition ed, 2010, Bosten u.a.: Pearson.

[9] Zimmer, M., H. Baars, and H. Kemper, *The Impact of Agility Requirements on Business Intelligence Architectures*, in System Science (HICSS), 2012 45th Hawaii International Conference on, 2012: IEEE.

[10] Thompson, W.J.J. and J.S. van der Walt, *Business intelligence in the cloud,* SA Journal of Information Management, 2010, Vol. 12 (1), p. 5 pages.

[11] Agarwal, S., et al., *Volley: Automated data placement for geo-distributed cloud services*, in Proceedings of the 7th USENIX conference on Networked systems design and implementation, 2010: USENIX Association.

[12] Kaviani, N., E. Wohlstadter, and R. Lea, *MANTICORE: a Framework for Partitioning Software Services for Hybrid Cloud*, in *4th ieee international conference on cloud computing technology and science*2012: Taipei, Taiwan.

[13] Ko, S.Y., K. Jeon, and R. Morales, *The hybrex model for confidentiality and privacy in cloud computing*, in Proceedings of the 2011 conference on Hot topics in Cloud Computing. USENIX Association, Portland, OR, 2011.

[14] Liu, C., B.T. Loo, and Y. Mao, *Declarative automated cloud resource orchestration*, in Proceedings of the 2nd ACM Symposium on Cloud Computing, 2011: ACM.

[15] Menzel, M. and R. Ranjan, *CloudGenius: Automated Decision Support for Migrating Multi-Component Enterprise Applications to Clouds*, in International World Wide Web Conference Committee (IW3C2), 2011.

[16] Menzel, M., et al., *(MC2) 2: A Generic Decision-Making Framework and its Application to Cloud Computing,* Software: Practice and Experience, 2011, Vol.

[17] Wieder, A., et al., *Orchestrating the Deployment of Computations in the Cloud with Conductor*, in Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation, 2012: USENIX Association.

[18] Chong, S., et al., *Building secure web applications with automatic partitioning,* Communications of the ACM, 2009, Vol. 52 (2), p. 79-87.

[19] Hunt, G.C. and M.L. Scott, *The Coign automatic distributed partitioning system,* Operating systems review, 1998, Vol. 33, p. 187-200.

[20] Apache-Foundation, *The Apache Hadoop project.,* 2013, Last Update, Date, Available from: http://hadoop.apache.org/.

[21] Baars, H. and H.G. Kemper, *Management Support with Structured and Unstructured Data – An Integrated Business Intelligence Framework,* Information Systems Management, 2008, Vol. 25 (2), p. 132-148.

[22] Shollo, A. and K. Kautz, *Towards an Understanding of Business Intelligence (Paper 86)*, in Australasian Conference on Information Systems 2010 (ACIS 2010), 2010, Brisbane (Australia).

[23] Negash, S., *Business Intelligence,* Communications of AIS, 2004, Vol. 13, p. 177-195.

[24] Golfarelli, M., S. Rizzi, and I. Cella, *Beyond data warehousing: what's next in business intelligence?*, in Proceedings of the 7th ACM international workshop on Data warehousing and OLAP, 2004: ACM.

[25] Marjanovic, O., *The Next Stage of Operational Business Intelligence: Creating New Challenges for Business Process Management*, in 40th Annual Hawaii International Conference on System Sciences (HICSS), 2007: IEEE.

[26] Manyika, J., et al., *Big data: The next frontier for innovation, competition, and productivity*, 2011.

[27] Jacobs, A., *The pathologies of big data,* Communications of the ACM, 2009, Vol. 52 (8), p. 36-44.

[28] Gutierrez, N., *Business Intelligence (BI) Governance,* Infosys White Paper, 2006.

[29] Baars, H. and H.G. Kemper, *Business Intelligence in the Cloud?*, in 14th Pacific Asia Conference on Information Systems (PACIS), 2010, Taipeh, Taiwan.

[30] Baars, H. and H.G. Kemper, *Ubiquitous Computing–an Application Domain for Business Intelligence in the Cloud?,* 2011, Vol.

[31] Thomson, W.J.J. and J.S. van der Walt, *Business intelligence in the cloud* South African Journal of Information Management, 2010, Vol. 12 (1).

[32] Qie, L. and H. Baars, *Die Cloud als neuer Ansatz zur Erhöhung der BI-Agilität?,* BI Spektrum, 2012, Vol. 7 (2), p. 26-29.

[33] Gartner, *Gartner Says Nearly One Third of Organizations Use or Plan to Use Cloud Offerings to Augment Business Intelligence Capabilities,* 2012, Last Update, Date [cited 2013 30th of January 2012], Available from: http://www.gartner.com/newsroom/id/1903814.

[34] Zeleny, M., *Multiple criteria decision making*, Vol. 25, 1982: McGraw-Hill New York.

[35] Saaty, T.L., *Theory and applications of analytic network process*, Vol. 4922, 2005: RWS publications Pittsburgh.

[36] Tran, V.X., H. Tsuji, and R. Masuda, *A new QoS ontology and its QoS-based ranking algorithm for Web services,* Simulation Modelling Practice and Theory, 2009, Vol. 17 (8), p. 1378-1398.

[37] Garg, S.K., S. Versteeg, and R. Buyya, *Smicloud: A framework for comparing and ranking cloud services*, in Utility and Cloud Computing (UCC), 2011 Fourth IEEE International Conference on, 2011: IEEE.