Decision Support for Application Migration to the Cloud
Challenges and Vision

Vasilios Andrikopoulos, Steve Strauch and Frank Leymann
Institute of Architecture of Application Systems, University of Stuttgart, Stuttgart, Germany

Keywords: Application Migration, Decision Support, Cloud-enabling Applications.

Abstract: The success of Cloud computing has encouraged many application developers to consider migrating their applications to the Cloud. Given the early market dominance of the IaaS service model, many existing works focus on selecting the best service provider for a set of criteria related to the virtualization and hosting of the application. In this work, we aim to progress the State of the Art by formulating a vision of a decision support system that incorporates multiple dimensions and different analysis tasks in feedback relationships with each other. The research challenges that need to be addressed towards this direction are identified and related to the different aspects of migration of applications to the various Cloud service models.

1 INTRODUCTION

The steadily increasing domination of Cloud computing in the software market means that existing applications may need to migrate to this environment in order to reap the benefits of reduced infrastructural costs and dynamic access to computational resources. While software has already started being developed specifically for the Cloud (forming Cloud-native applications), existing systems must be adapted to be suitable for the Cloud, requiring to make them Cloud-enabled (Andrikopoulos et al., 2013). Existing work, e.g. (Li et al., 2010; Khajeh-Hosseini et al., 2012; Menzel and Ranjan, 2012), focuses on the IaaS service model (Mell and Grance, 2011), building on the success and early market dominance of infrastructure virtualization solutions like Amazon Web Services. In this context, providing decision support for the migration of applications to the Cloud is reduced to the selection of the appropriate IaaS provider that can support the best application performance for the least cost.

Beyond this VM-based migration of applications to the Cloud however, new PaaS and SaaS offerings of Cloud providers enable alternative migration options for applications. Data stores like the Amazon Relational Database Service (Amazon RDS), and business logic executing environments like the Google App Engine allow for partially or fully migrating the application. Decision support in this case needs also to include different options for segmenting and distributing the application in the Cloud, figuring out which parts of the application should scale, and how multi-tenancy can be implemented to maximize the return of investment for application developers.

The position that this work takes is that migrating to the Cloud is a multi-dimensional problem with multiple decision points that may create feedback loops with each other, and with various analysis tasks related to them. The contribution of this paper is to identify these decision points and analysis tasks, and define their dependencies. Based on these results, a vision of a Cloud migration decision support system considering these aspects is provided. In order to reach this point however, a comprehensive analysis of the challenges related to migrating to the Cloud is required.

The rest of this paper is structured as follows: Section 2 summarizes the State of the Art in decision support systems for Cloud migration. Section 3 discusses the issues in migrating applications to the Cloud and identifies the research challenges for supporting the decisions related to them. Section 4 presents our vision for how a decision support system can address these challenges; Section 5 concludes this paper.

2 STATE OF THE ART

Existing work on decision support for application migration to the Cloud identify a number of criteria to be considered, with the cost being the most impor-

Andrikopoulos V., Strauch S. and Leymann F..
Decision Support for Application Migration to the Cloud - Challenges and Vision.
DOI: 10.5220/0004407001490155
In *Proceedings of the 3rd International Conference on Cloud Computing and Services Science* (CLOSER-2013), pages 149-155
ISBN: 978-989-8565-52-5

tant one. The work presented in (Khajeh-Hosseini et al., 2011) discusses two decision support tools implemented as spreadsheets: a cost calculation one, and a risk-benefit analysis one. The former is geared towards IaaS solutions, and allows users to calculate the cost for their infrastructure when virtualized and ran in one or more VMs on the Cloud. The latter surveys and categorizes a number of benefits and risks identified in the literature (also for IaaS solutions), and informs about potential consequences, mitigation approaches and indicators for each risk. The two tools are then combined in the Cloud Adoption Toolkit (Khajeh-Hosseini et al., 2012) that supports decisions on technology suitability, cost, energy consumption, stakeholder impact and operational feasibility.

The $(MC^2)^2$ framework (Menzel et al., 2011), later incorporated in the CloudGenius toolkit (Menzel and Ranjan, 2012), also focuses on IaaS solutions and provides a multi-criteria approach in decision support for application migration. Both approaches allow users to define multiple quantitative and qualitative requirements that are then matched against a knowledge base of Cloud service providers. Various criteria to evaluate the possible alternatives are identified, and a series of cost categories are considered. Beyond the emphasis on the IaaS service model, the main criticism of all of the above works is the lack of connecting the outcomes of each analysis step with the particular decisions to be made. In this sense, system users are able to evaluate their design, but they do not have any insight into how each design decision affects the overall system.

In particular with respect to selecting appropriate Cloud service providers for the migration, Brebner and Liu (Brebner and Liu, 2011) report on the costs of running applications with different work loads in various Amazon, Google and Microsoft offerings. A similar calculation is reported in (Li et al., 2010) and (Suleiman et al., 2011) for anonymized Cloud services. The benchmarking of (Li et al., 2010) takes into account storage and networking costs, while the analysis of (Suleiman et al., 2011) focuses on the elasticity options of the offerings. While these works can be used as part of the decision support process, they do not consider the characteristics of the application itself it terms of how it needs to be adapted to operate in the Cloud environment.

3 CLOUD MIGRATION CHALLENGES

As the first step towards a decision support system beyond the State of the Art, in the following we use the issues and concerns identified in previous work on application migration (Andrikopoulos et al., 2013) as the basis to identify the research challenges affecting decision support for Cloud migration.

Application Distribution. Logically and physically distributing the data and the computational aspect of applications over the Cloud creates a series of economic, performance, and legal issues. The calculations in (Armbrust et al., 2009) illustrate the trend of computational power becoming cheaper faster than disk storage, which becomes cheaper faster than network bandwidth in turn. It is therefore in principle good practice to keep data and computation close to each other. The actual physical *proximity* of the migrated application components and their data is however beyond the control of the application developer — except from the location of the Cloud provider data center.

Similarly, it becomes very difficult to ensure *compliance* to legal and regulatory requirements, concerning for example the privacy of users. EU regulations for instance require the physical location of the stored data to be inside the EU borders (Reese, 2009). Even if it is possible to define the region for deploying the application, the cost of using them can vary due to different charging policies (Suleiman et al., 2011). Furthermore, previous work has shown that Cloud services performance varies significantly for different regions (Schad et al., 2010), or even for different data centers inside the same region (Li et al., 2010).

It is important to notice that distributing the application may be necessary in order to address the need for business resiliency (Andrikopoulos et al., 2013), while aiming to minimize the risk (both in terms of security and QoS assurances) of the application migration to the Cloud. This requires backup solutions and use of multiple Cloud providers as alternatives (see also (Badger et al., 2012)), translates into a multiplication of the costs, and may require application re-design.

Research Challenges. *How to estimate the effect of distributing logically and physically the application in terms of performance and cost? How to address compliance across different jurisdictional domains? What is the trade-off between business resiliency and additional incurred costs?*

Elasticity. Elasticity provides the means for optimizing resource usage in case of fluctuating and/or unknown application work loads. Ideally, the consumer perceives an infinite number of resources available in any quantity, at any time. Existing works like (Brebner, 2012; Suleiman et al., 2011; Vaquero et al., 2011) however connect the benefits from elasticity mechanisms not only with the Cloud solutions themselves, but also with the particular characteristics of the application in terms of its work load. Two types of scalability are identified in (Vaquero et al., 2011): *horizontal*, i.e. addition of more application instances where required, and *vertical*, that is, addition of more computational resources to the application. While vertical scalability is possible in principle for all applications, it largely depends on the service provider to offer the mechanisms to implement scaling dynamically. Horizontal scalability on the other hand mostly depends on the application components and the application as a whole to support it as an option. Scaling a stateful component like a database, for example, requires additionally the choice of an appropriate data replication mechanism (Andrikopoulos et al., 2013).

Different elasticity strategies can be implemented depending on the decisions to be made on *a*) *what* part of the application to scale, *b*) *how much* it is necessary to scale taking also into account issues related to licensing costs for multiple VMs, *c*) *which scaling type* to use and, *d*) *when* and *how fast* is it necessary to scale, considering the observed effect of scaling latency (Brebner, 2012) for starting new images in the Cloud. Furthermore, the choice of the Cloud deployment model also imposes limitations to the benefits of elasticity, as identified by (Badger et al., 2012). Private Cloud solutions, for example, impose essentially the same limitations in maximum capacity as in the case of traditional data centers.

Research Challenges. *How to decide which elasticity strategy to implement for a given application? What is the effect of this decision to the operational costs of the application under different work loads?*

Multi-tenancy. The multi-tenant model of serving multiple consumers from a common pool of computing resources, including storage, processing, memory and network bandwidth, is one of the essential characteristics of Cloud computing (Mell and Grance, 2011). There are two fundamental aspects of multi-tenancy awareness (Strauch et al., 2012a): *communication*, i.e. supporting message exchanges isolated per tenant, and *administration and management*, i.e. allowing each tenant to configure and manage individually their communication endpoints at application or service level. Tenant isolation is further decomposed into data and performance isolation between tenants of the same system. Existing approaches on enabling multi-tenancy typically focus on different types of isolation in multi-tenant applications for the SaaS delivery model, see for example (Guo et al., 2007); most of the work however, especially for performance isolation, remains an open research challenge (see for example (Alexandrov et al., 2012; Krebs et al., 2012)).

Furthermore, as discussed in (Andrikopoulos et al., 2013), the three Cloud service models differ significantly in the granularity of the functionality provided to the consumer, and the required capability of the consumer to manage and control the underlying Cloud infrastructure. The result is a trade-off between the effort required by the consumer to enable multi-tenancy on application level (especially in the IaaS model), and the responsibility of the provider to offer the suitable mechanisms for this purpose (as exemplified in the SaaS model).

Research Challenges. *Which Cloud service model (IaaS, PaaS, SaaS) to choose when considering the multi-tenancy needs of the application? What is the effort required to implement multi-tenancy for the chosen service model?*

QoS. Migrating an application to the Cloud entails a loss of control over the QoS characteristics of the application due to the reliance on the QoS levels offered by the service provider. As a result, the QoS characteristics offered by a Cloud service provider appear to have a greater importance to application stakeholders than hosting the application in a non-Cloud environment. The most obvious case of this dependence is the effect of service provider outages to businesses hosting their applications with them. As discussed in (Badger et al., 2012), availability is usually calculated in practice on the basis of the billing cycle, which may vary from minutes to hours, and not on the total up-time of the service. Furthermore, failure to provide this availability is typically compensated with service credit in future use of the services. In any case, the overall availability of a migrated application can actually vary significantly with respect to the availability of the used services, depending on the actual behavior of the application. Safely predicting the availability of such an application is largely an open issue.

Beyond availability, application QoS levels in the Cloud are also affected by two major factors: the *performance variability* of the Cloud providers, and the *network latency* between the Cloud service consumers and the service (Andrikopoulos et al., 2013). While performance variability may be (partially) addressed by benchmarking Cloud service providers to select an appropriate one in the manner of (Li et al., 2010), network latency is in the general case beyond the control of the application provider. SLAs with existing application consumers are therefore difficult to be enforced in this environment.

Research Challenges. *What is the acceptable level of Cloud service availability and performance variability for an application with respect to existing and future SLAs? How to factor and address network latency in the performance analysis? How does the choice of service provider affects this analysis?*

Migration & Operation Cost. As discussed in Section 2, the focus of existing work on decision support systems is on selecting the best provider in terms of costs for a given application (Hajjat et al., 2010; Khajeh-Hosseini et al., 2012; Menzel and Ranjan, 2012), and under different work loads (Li et al., 2010; Brebner and Liu, 2011). Beyond provider selection however, cost analysis should also consider a number of other aspects. Firstly, different providers offer different pricing models, see for example (Suleiman et al., 2011), and combinations of these models, e.g. subscription to a service like Windows Azure offering a discount on the per-use charging for given usage quotas, beyond which regular per-use rates apply. More complicated models have also starting arising, with Amazon Web Services for example allowing customers to bid for unused EC2 capacity by means of Spot Instances (Amazon Web Services, 2012).

Furthermore, software licensing is often based on the number of CPUs which does not fit the dynamicity of the solutions offered by the Cloud (Reese, 2009). For example, scaling a system may easily result in unintended license agreement violations. In practice, some Cloud providers offer different licensing options to their consumers. In Amazon RDS for example, consumers can bring their own license for MySQL, Oracle or IBM DB2, get charged for a per-hour license using Oracle DB, or pay a one-time charge per Amazon RDS instance to get reduced hourly charging rates (Suleiman et al., 2011). Some companies include licensing fees for free with each account.

The cost analysis for the migration of an application to the Cloud requires also to compare the costs of operating in the Cloud versus operating in a local data center. The discussion in (Armbrust et al., 2009) illustrates the factors that allow Cloud computing to be more profitable than on-premises hosting: pay per use only, avoid data center power, cooling and building costs, out-sourcing of the operational costs to the provider. Existing works identify a relation between the type and work load of the application, and the benefits by migrating the operation to the Cloud. In (Tak et al., 2011) for example, the authors conclude that complete application migration is beneficiary only for small or stable organizations, while partially migrating some of the application components to the Cloud is too expensive due to the high costs of data transfer. Reese posits that, in principle, significant cost savings due to the transfer from capital expenses (CAPEX) to operating expenses (OPEX) incur only as the variance increases between peak, average and low capacity of the system (Reese, 2009).

Finally, cost analysis should additionally consider the effort required for the adaptation of the application in order to operate in the Cloud. In (Andrikopoulos et al., 2013), the authors identify four types of application migration to the Cloud: *replacement* of existing application components with Cloud offerings, *partial migration* of some of the application functionality to the Cloud, *migration of the complete software stack*, and *cloudification* of the application by replacing all of its components with Cloud services. Depending on the type of the application, and the architectural layer that is affected by the migration, the authors conclude that different degrees of adaptation effort may be required.

Research Challenges. *How to decide on the best provider in terms of costs in the presence of different pricing models? What is the trade-off between cost and performance? How to decide which deployment option (completely locally, fully on the Cloud, combination of the above) fits better the application work load profile in question? What is the additional effort required for adapting the application to operate in the Cloud?*

Security & Data Confidentiality. Security and confidentiality of data are major concerns and obstacles for many enterprises to migrate to the Cloud (Armbrust et al., 2009) and for this reason they need to be addressed separately from the other QoS concerns. Security entails both the *communication* and *data* aspects, but also the *physical/digital* one, i.e. the risk of losing or compromising data due to data center failures or other physical attacks. In principle, security is a shared concern between Cloud service providers and consumers (Andrikopoulos et al., 2013). On one hand, service providers are required to offer security-enabling mechanisms like encrypted communications. On the other hand however, application developers are the ones responsible for adapting the migrated application accordingly to use these mechanisms and further configure it appropriately. Ultimately, it is also the responsibility of the application users to interact with it in a secure manner following for example the security recommendations discussed in (Badger et al., 2012).

In addition to security, confidentiality of Cloud-migrated data is an important concern for all kinds of organizations. Confidentiality refers to keeping critical data private and secure. Critical data may be business secrets, personal data, health care data and so on. (Strauch et al., 2012b; Andrikopoulos et al., 2013) provide a number of techniques like pseudonymization and anonymization that can be used to ensure the confidentiality of migrated data. However as they point out, these techniques are not sufficient to avoid data breaches due to e.g. attacks from inside the Cloud environment, and the leaking of information while processing these data in the Cloud. Data retention policies beyond contract terms are also not explicitly defined for most of the existing providers, creating possible sources of data breaches.

Research Challenges: *What is the effort to (re-)engineer the application to incorporate security mechanisms like encryption of data? Which part of the application data can be characterized as critical, and under which conditions can it be migrated to the Cloud securely?*

4 DECISION SUPPORT FOR CLOUD MIGRATION

What becomes obvious from the previous discussion is that migrating an application to the Cloud requires making a number of decisions related to the how the application is to be refactored for the Cloud, how it is supposed to scale, an so on. In this sense, cost analysis and performance prediction, which has been the focus of the SotA as discussed in Section 2 provide only partial support for these purposes. Addressing these concerns, Fig. 1 provides an overview of the concepts required to implement a decision support system to support the migration of applications to the Cloud. This conceptual view defines our vision for what constitutes a complete solution for application developers and stakeholders alike that are considering whether and how to migrate their application to the Cloud.

More specifically, two types of concepts are identified: *decisions* that need to be made (and therefore are the focus of the system), and *tasks* that need to be performed in order to support these decisions — and that in turn affect their outcome. The key decision that the system supports refer to:

1. *How to distribute the application* across service providers, and between the Cloud and the local data center. Different application topologies based on the various identified migration types (Andrikopoulos et al., 2013) may be considered.
2. *How to select a (Cloud) service provider and offering* that fits the application needs in terms of cost, expected performance, compliance requirements and security concerns.
3. *What are the requirements of the application in terms of multi-tenancy*, i.e. to what extent the existing application is required to support multi-tenancy, to what degree it is designed for this purpose, and how it should be (re-)engineered to support multi-tenancy.
4. *Which is the elasticity strategy* that the application needs to implement in order to cope with its demand in the face of SLAs and expectations of its users.

Each of these decisions has a direct or unintentional influence on the others, illustrated with transparent arrows in Fig. 1. Deciding on the appropriate elasticity strategy for example is meaningless without the Cloud service provider supporting it. The selection of service provider includes the decision which service delivery model to use. As a result, both the options in distributing the application and how multi-tenancy can be implemented is directly influenced by it.

In addition to the decisions and their relationships, the following tasks are also identified:

- **Work Load Profiling**: Defining or estimating the expected work load profile of the application is prerequisite for both other tasks to be performed (performance calculation and cost analysis), as well as input for any decision on how to distribute the application, and which elasticity strategy fits this profile.
- **Compliance Assurance.** Ensuring the compliance to regulations regarding, e.g., privacy of personal data, affects directly both the selection of the provider (especially in terms of location of the data service), as well as how the application is to be distributed so that e.g. personal data may need to be retained on premises.
- **Identification of Security Concerns.** Defining which data and communications are critical to be protected drives the selection of an appropriate service provider that fulfills these security constraints. Delivering the application in a multi-tenant manner imposes further constraints with respect to data isolation that need to be considered for this task.
- **Identification of Acceptable QoS Levels.** Based on existing and planned SLAs, acceptable levels for QoS characteristics like availability of the service provider can be inferred. Beyond service provider selection, this task also drives the definition of an appropriate elasticity strategy to ensure

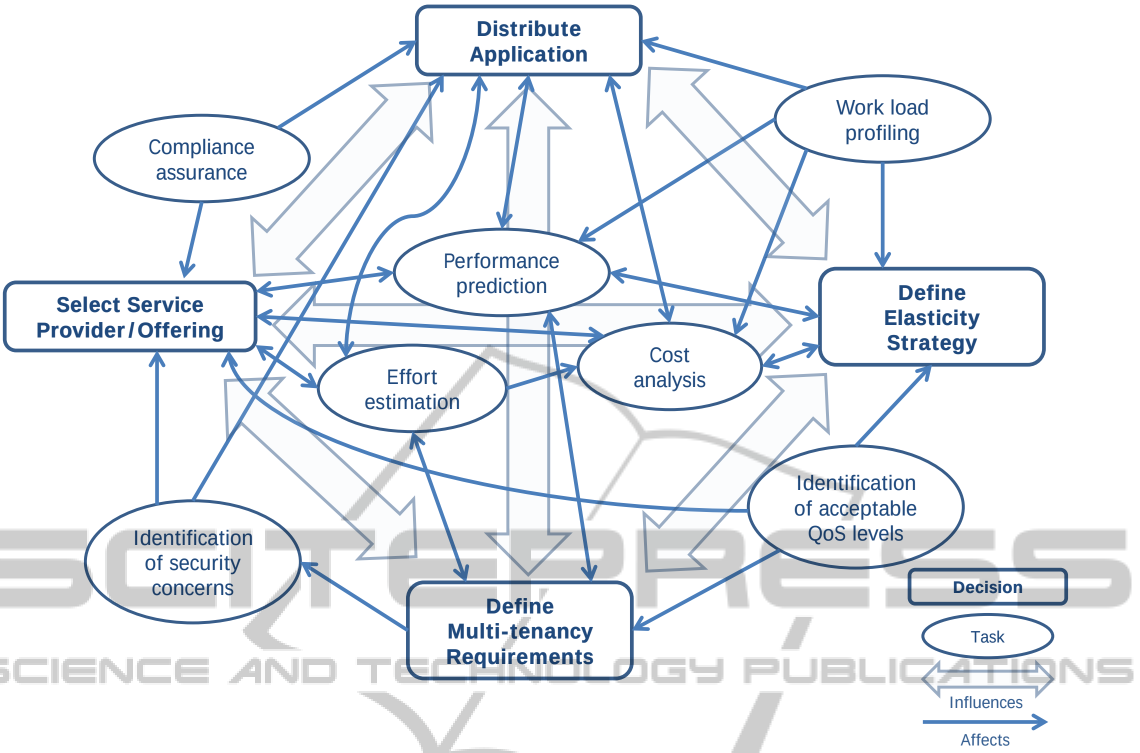

Figure 1: Decision Support System for Cloud Migration: conceptual view.

these QoS levels (for the given work load profile) and constrains the options for application multi-tenancy.

- **Performance Prediction.** Depending on how the application is distributed across different service providers, which providers where selected for this purpose, and which elasticity strategy was defined, this task creates projections about the non-functional behavior of the application after it is migrated to the Cloud. These projections can in turn be used to change any of the previous decisions, starting a new feedback loop between them.
- **Cost Analysis.** Similarly to performance prediction, cost analysis both affects and is affected by the application distribution, elasticity strategy and service provider selection decisions. Furthermore however, this task should also consider the effort estimation for adapting the application to operate in the Cloud.
- **Effort Estimation.** This task focuses on providing an estimate related to the amount of work required to adapt the application depending on the type of migration that is considered, and the architectural layer that is affected. As a result, this task requires input from the application distribution, service provider selection and multi-tenancy requirements decisions, and may result in changes to them in turn.

As shown in Fig. 1, decisions and tasks form a *network* of relationships, providing multiple points of entry to the decision making process of migrating an application to the Cloud. A Cloud migration decision support system should therefore support this process not in a structured hierarchical manner, but rather in a networked fashion, allowing application designers to provide different answers for each decision that needs to be made, ask them for the appropriate input required to perform the identified tasks, and demonstrate the effects of their decisions to concrete goals like the cost and performance of the application. Such an approach is missing from the current State of the Art and constitutes our future work in the area.

5 CONCLUSIONS

The position we took in this work is that supporting the decision making process in migrating the application to the Cloud is a multi-dimensional problem with multiple decisions to be taken: how to segment and distribute the application, which Cloud service provider and offering to choose, which elasticity strategy fits better the application, and what are the requirements in implementing application multi-tenancy. These decisions influence each other, and depend on a number of tasks like cost analysis and performance prediction that need to be performed in order to support them. Related to these decisions and tasks are a series of research challenges stemming from the need to adapt the application to operate in the Cloud environment that also need to be addressed. Our vision for a decision support system for migration of applications to the Cloud as discussed here considers all of the above points in order to advance the State of the Art in the field.

ACKNOWLEDGEMENTS

The research leading to these results has received funding from the 4CaaSt project (http://www.4caast.eu) part of the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 258862.

REFERENCES

Alexandrov, A., Folkerts, E., Sachs, K., Iosup, A., Markl, V., and Tosun, C. (2012). Benchmarking in the Cloud: What it Should, Can, and Cannot Be. In *4th TPC Technology Conference on Performance Evaluation & Benchmarking (TPCTC), VLDB 2012*.

Amazon Web Services (2012). How AWS Pricing Works. http://aws.amazon.com/whitepapers/.

Andrikopoulos, V., Binz, T., Leymann, F., and Strauch, S. (2013). How to adapt applications for the cloud environment. *Computing (to appear)*. http://dx.doi.org/10.1007/s00607-012-0248-2.

Armbrust, M. et al. (2009). Above the Clouds: A Berkeley View of Cloud Computing. Technical Report UCB/EECS-2009-28, EECS Department, University of California, Berkeley.

Badger, L., Grance, T., R., P.-C., and Voas, J. (2012). Cloud computing synopsis and recommendations - recommendations of the national institute of standards and technology. NIST Special Publication 800-146.

Brebner, P. (2012). Is your cloud elastic enough?: performance modelling the elasticity of infrastructure as a service (iaas) cloud applications. In *Proceedings of ICPE'12*, pages 263–266. ACM.

Brebner, P. and Liu, A. (2011). Performance and cost assessment of cloud services. *Service-Oriented Computing*, pages 39–50.

Guo, C., Sun, W., Huang, Y., Wang, Z., and Gao, B. (2007). A Framework for Native Multi-Tenancy Application Development and Management. In *Proceedings of CEC/EEE'07*, pages 551–558. IEEE.

Hajjat, M., Sun, X., Sung, Y., Maltz, D., Rao, S., Sripanidkulchai, K., and Tawarmalani, M. (2010). Cloudward bound: planning for beneficial migration of enterprise applications to the cloud. In *ACM SIGCOMM Computer Communication Review*, volume 40, pages 243–254. ACM.

Khajeh-Hosseini, A., Greenwood, D., Smith, J. W., and Sommerville, I. (2012). The cloud adoption toolkit: supporting cloud adoption decisions in the enterprise. *Software: Practice and Experience*, 42(4):447–465.

Khajeh-Hosseini, A., Sommerville, I., Bogaerts, J., and Teregowda, P. (2011). Decision support tools for cloud migration in the enterprise. In *2011 IEEE International Conference on Cloud Computing (CLOUD)*, pages 541–548. IEEE.

Krebs, R., Momm, C., and Kounev, S. (2012). Metrics and Techniques for Quantifying Performance Isolation in Cloud Environments. In Buhnova, B. and Vallecillo, A., editors, *Proceedings of the 8th International ACM SIGSOFT Conference on the Quality of Software Architectures, CBSE'12*, pages 91–100, New York, USA. ACM Press.

Li, A., Yang, X., Kandula, S., and Zhang, M. (2010). Cloudcmp: comparing public cloud providers. In *Proceedings of the 10th annual conference on Internet measurement*, IMC '10, pages 1–14, New York, NY, USA. ACM.

Mell, P. and Grance, T. (2011). The nist definition of cloud computing (draft). *NIST special publication*, 800:145.

Menzel, M. and Ranjan, R. (2012). Cloudgenius: decision support for web server cloud migration. In *Proceedings of WWW '12*, pages 979–988, New York, NY, USA. ACM.

Menzel, M., Schönherr, M., and Tai, S. (2011). (mc2) 2: criteria, requirements and a software prototype for cloud infrastructure decisions. *Software: Practice and Experience*.

Reese, G. (2009). *Cloud application architectures*. O'Reilly Media, Inc.

Schad, J., Dittrich, J., and Quiané-Ruiz, J. (2010). Runtime measurements in the cloud: observing, analyzing, and reducing variance. *Proceedings of the VLDB Endowment*, 3(1-2):460–471.

Strauch, S., Andrikopoulos, V., Gómez Sáez, S., Leymann, F., and Muhler, D. (2012a). Enabling Tenant-Aware Administration and Management for JBI Environments. In *Proceedings of SOCA'12*. IEEE Computer Society Conference Publishing Services.

Strauch, S., Breitenbücher, U., Kopp, O., Leymann, F., and Unger, T. (2012b). Cloud Data Patterns for Confidentiality. In *Proceedings of CLOSER'12*, pages 387–394. SciTePress.

Suleiman, B., Sakr, S., Jeffery, R., and Liu, A. (2011). On understanding the economics and elasticity challenges of deploying business applications on public cloud infrastructure. *Journal of Internet Services and Applications*, pages 1–21.

Tak, B., Urgaonkar, B., and Sivasubramaniam, A. (2011). To move or not to move: The economics of cloud computing. In *Third USENIX Workshop on Hot Topics in Cloud Computing (HOTCLOUD 2011)*.

Vaquero, L., Rodero-Merino, L., and Buyya, R. (2011). Dynamically scaling applications in the cloud. *ACM SIGCOMM Computer Communication Review*, 41(1):45–52.

All links were last followed on March 12, 2013.