Cloud Service Recommendation and Selection for Enterprises

Shixing Yan*, Chunqing Chen*, Guopeng Zhao*
* HP Labs Singapore
1 Fusionopolis Way
Singapore 138632
Email: {shixing.yan, chunqing.chen, guopeng.zhao}@hp.com

Bu Sung Lee*†
† School of Computer Engineering
Nanyang Technological University
Singapore 639798
Email: ebslee@ntu.edu.sg

***Abstract*—The standardization of cloud services makes it possible to have one cloud service management platform for customers to take advantage of their subscriptions from various cloud providers. One immediate benefit lies in the way that enterprise customers may have plenty of choices of available providers, when seeking cloud services to fulfill their criteria, such as price or Service Level Agreement (SLA). However, how to select appropriate cloud offerings in terms of business requirements and company policies becomes challenging and non-trivial, especially when a composition of multiple services are to be chosen for achieving business goals. In this paper, a systematic framework on top of a hybrid cloud management platform is proposed for enterprises to automatically recommend and select cloud services according to business requirements, company policies and standards, and the specifications of cloud offerings.**

***Keywords*—Cloud Service Selection; Cloud Management; Hybrid Cloud; MCDM**

I. Introduction

Cloud computing [1], [2] has been considered by more and more IT decision makers in public sector and large enterprises to support their organizations business. We have seen the trend of accelerated cloud adoption because the enterprises are now seeing cloud computing as the key element of the IT arsenal with the promise of faster time to market, financial savings and immense scalability on pay-per-use mode. Although public cloud services offer the customer great business agility with no upfront investment, the public sector and large enterprises are currently focused on private and hybrid clouds because of the concerns on security, integration with existing system, and regulatory/compliance issues. With the trend of moving to hybrid cloud and leveraging the benefits brought by cloud computing, it brings up the manageability issue of the cloud infrastructure in hybrid cloud environment due to lack of standards for configuration and management of cloud services across cloud providers. Several approaches [3]–[6] have been proposed recently to manage heterogeneous cloud services from different providers. Several international standardization organizations are also working on the cloud standards to increase the manageability and interoperability of cloud services, such as DMTF's Cloud Management Work Group (CMWG) working on the Cloud Infrastructure Management Interface (CIMI) [7]. These effort will greatly help the enterprises embrace the cloud to replace their traditional IT in the long run.

We have seen the rapid growth of the number of cloud services available in this emerging market as service providers have seen the cloud as a catalyst for their revenue growth. This provides users a number of cloud services range from infrastructure-as-a-service (IaaS), through platform-as-a-service (PaaS), to software-as-a-service (SaaS). The cloud services market is being commoditized with the big variety of cloud offerings. Many cloud providers have rolled out "App store"-like cloud service marketplace to allow their customers to find, compare and subscribe cloud services in their clouds, e.g., AWS Marketplace [8], Saleforce.com's appexchange [9], VMware's Solution Exchange [10], The Google Apps Marketplace [11]. There are also some websites, HP Cloud Service Catalog [12], CloudSurfing [13], DreamSimplicity [14], which provide a catalog of cloud services from various cloud providers. Enterprises will be able to achieve their business goals by speeding their innovations and improving their operational efficiency through adopting the appropriate cloud services, in response to their business needs. However, there are tedious business procedures for traditional IT procurement in the enterprises to guarantee the IT services/applications complying with company policies and standards, and now it becomes a major obstacle for enterprises to select the cloud services efficiently in response to immediate business needs.

We have developed a policy-based cloud management system called *Monsoon* which enables enterprises to deploy business workloads in hybrid cloud and manage the full life-cycle of their cloud infrastructures [5]. Monsoon provides a multi-tenant self-service portal for enterprises to build, package, and provision cloud infrastructure services for their internal users through a unified IaaS proxy service [15]. The system has evolved to manage more generic cloud services from managing cloud infrastructure services only, and this makes it possible to build an enterprise cloud service catalog to allow the internal users consuming cloud services from various cloud providers across both private and public clouds. There is also a cloud sourcing management solution for enterprise

978-3-901882-48-7 ©2012 IFIP

IT, HP Cloud Service Catalog [12], to effectively discover SaaS and manage their cloud service portfolio in public cloud. These works have laid the foundation for us with providing the enterprise customers a cloud service catalog with a vast variety of cloud services available in the platform, and the cloud services may come from different cloud providers with diverse features and pricing schemes. How to select the most appropriate cloud services according to the business requirements and company policy becomes a new challenge for enterprises when the cloud service catalog is available. There have been some existing work on cloud service selection [16]–[22], but they did not provide a full solution to address the specific requirement of cloud service selection for enterprises, especially on supporting complex interdependent relations and company policy compliance. The existing work on cloud service selection mostly focus on evaluation and assessment of available cloud services according to user requirements, while the other missing part is that they cannot help users refine their requirements when there is no feasible solution.

In this paper, we present a framework which aims to provide enterprises a cloud service recommendation and selection tool to accelerate and automate the cloud adoption to meet the dynamic business requirements and comply with the company policies and standards. We discuss the generic framework and our most current research results, and some research work in progress are also shared in this paper. In our approach, we firstly address the policy compliance issue and provide automatic conflict detection and explanation by Policy Implementation Engine (PIE) which was briefly introduced for the first time in [5]. The technical details of PIE have been discussed in our recently released technical report [23]. PIE will generate a group of feasible solutions or help user refine their requirements if there is no feasible solution. The Cloud Intelligence module will optimize the cloud service selections from PIE according to the user preference on different criteria ranging from quantitative metrics to collective intelligence from social media.

The reminder of the paper is organized as follows. Section II gives a brief introduction on cloud services, utility function and multiple criteria decision making. We elaborate the design of our framework for cloud service recommendation and selection in Section III. Section IV concludes the paper.

II. Background

A. Cloud Services

We have characterized the cloud services in different features in order to build up a framework to recommend and select cloud services for enterprises. For example, some features of a typical cloud services include service availability, price, and the geographical location of the data center, etc. Some cloud services provide similar features and therefore they could be the options for satisfying a particular customer requirement, for example, IaaS services usually consist of the features like processor, memory, network, and storage. Our early work on managing heterogeneous cloud services from multiple providers [5], [15] enables enterprises to distribute their work loads into various cloud services to compose a complicated services.

For example, when we consider deploying a three-tier application into cloud, we may want to deploy the virtual machine (VM) at the application layer and the VM at the data layer in the same subnet to minimize network latency. We also require another VM which is the backup database server to be placed in another location (which is different than the first two VMs) for the sake of fault-tolerance. This example shows up another research challenge on cloud service selection brought by the interdependent relations among the user requirements. Expressive interdependent relations can increase the difficulties of finding appropriate services since one user requirement may affect or be affected by others. Furthermore, the interdependencies can easily introduce subtle conflicts which cause these requirements to be impossible to satisfy.

B. Utility Function

Utility function has been widely used by economist to measure consumer's preferences over some set of goods and services. In our approach, we see that cloud services available in the cloud service catalog are considered as commodities, and therefore we can construct the user's utility function to represent his or her preferences over certain set of cloud services according to the requirements. Since the user evaluates a number of alternatives on the basis of two or more criteria or attributes, the cloud service's common features need to be modeled, quantified, and then normalized to be represented in utility function. The higher the value of utility function represents the greater benefit that users gain from the composition of cloud services. Hence, the maximization of the utility function with given constraints will be the ideal outcome of the service recommendation and selection.

C. Multiple Criteria Decision Making

Multiple Criteria Decision Making (MCDM) is a systematic procedure that the decision maker chooses among a number of alternatives while he or she needs to evaluate on the basis of two or more criteria. This research field has been widely studied as MCDM is practical in our daily lives or in professional settings. Many different techniques have been developed to tackle the MCDM research challenge [24], [25], such as goal programming, vector optimization algorithms, and interactive approaches, etc.

III. Framework Architecture

We propose a framework to help enterprises make the choices of cloud service subscriptions based on the the policy-based service selection and rating-based service recommendations according to user requirements and the specifications of cloud offerings in the platform. The conceptual architecture of this new framework is shown in Figure 1. We will discuss different components and modules of the framework in the following sub-sections.

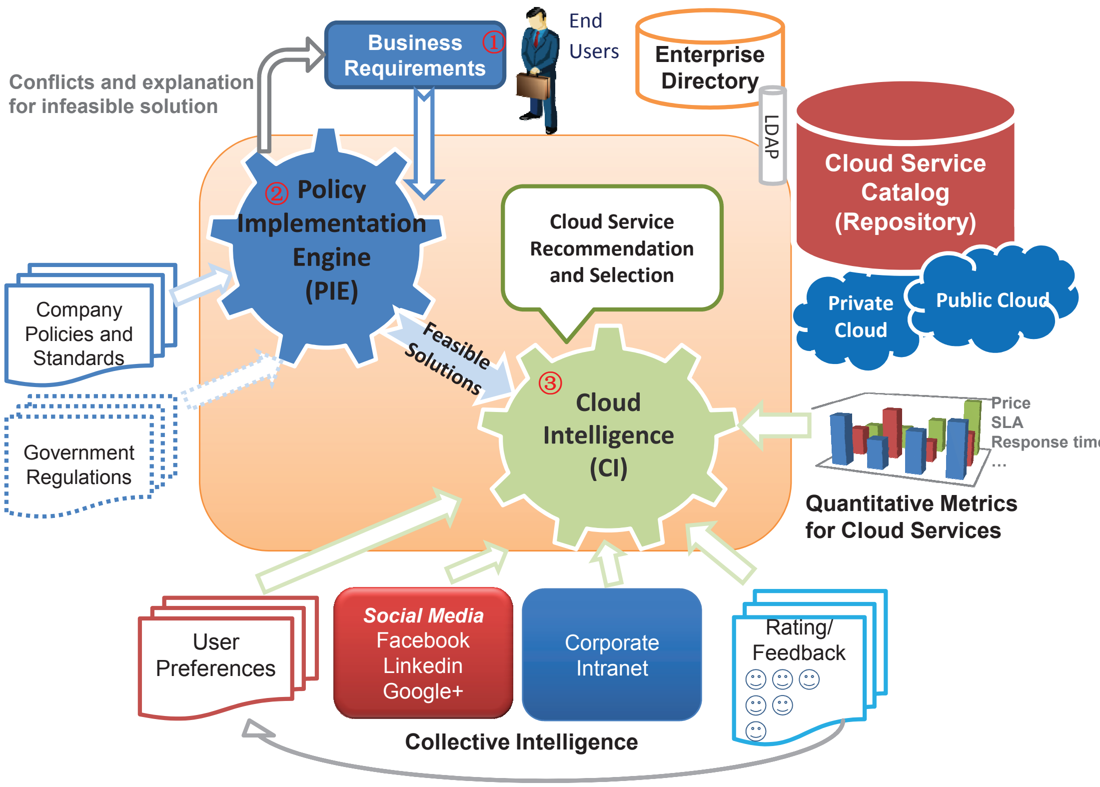

Fig. 1. Framework of cloud service recommendation and selection for enterprises.

A. *Cloud Service Modeling and Feature Extraction*

The diversity of cloud services makes the comparison of cloud offerings difficult, though many of them have common features and can be the alternatives to some other services. We firstly classify the cloud services to the catalog, and extract the categorized features to facilitate the analysis for optimizing the cloud service selection. As we can see from Figure 1, the cloud service catalog serves as a repository to store all the specifications of various cloud service offerings from cloud providers which can be in either private or public cloud. Cloud service catalog is provided in the framework for enterprises to discover and procure a wide variety of cloud services to satisfy their business needs. The hybrid cloud management solution, Monsoon, is able to embed the cloud service catalog to provide the full life-cycle service management for enterprises in this framework [5].

B. *Quantitative Metrics of Cloud Services*

The commonality of the features extracted from cloud services enables the measurement in the comparisons among cloud services. The business requirements from users can be translated to a set of criteria based on the cloud services' common features. The utility function of a set of cloud services will be represented by a performance score on the criteria derived from the user requirements. For example, a customer requires a group of VMs with relevant information: the type of supported OS, locations of the service, and the price and OS with the RAM size. In our approach, we firstly quantify these service features to make each service measurable on the criteria. The quantified performance scores are further normalized to facilitate the comparison among the criteria.

C. *Policy Implementation Engine*

Major concern for enterprises wishing to adopt cloud computing is compliance with IT policies mandated by the enterprise or government [26], e.g., data sovereignty issues in cloud computing. Unfortunately, it is troublesome and impractical to manually detect various conflicts from a number of requirements, and even harder to identify problematic requirements that cause conflicts. It is hence important to provide enterprise users with comprehensive support which enables conflict detection and explanation for their cloud service selection. Policy Implementation Engine (PIE) is designed to consider interdependencies that span across user requirements, and also the compliance of his/her requirements with relevant enterprise policies and government regulations. PIE is a core component in the framework as shown in Figure 1. The preliminary design of PIE has been discussed in [5]. Recently, we have improved PIE based on formal verification and constraint solving techniques, and PIE has been implemented in Monsoon for enterprises to automatically checks various conflicts covering the violation against company policies and government regulations, and inconsistency within user requirements [23]. The latest design of PIE is shown in Figure 2 [23].

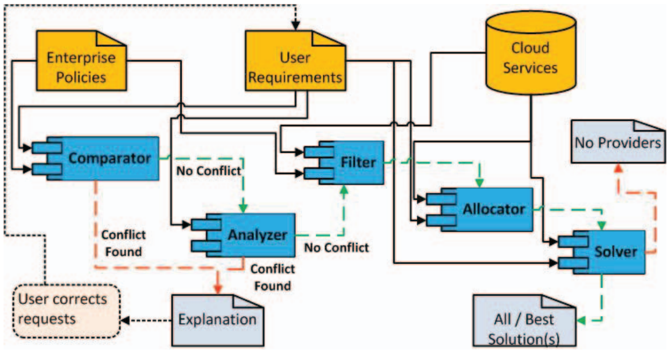

Fig. 2. Policy Implementation Engine (PIE)

PIE firstly checks the existence of conflicts based on user requirements as shown in Figure 1. The conflicts detected by PIE are categorized into two types based on the causes, and then they are prioritized according to their impact scopes. When a conflict is detected, a form containing an explanation of the conflict is generated to assist users to quickly understand and resolve the conflict. This makes the process of analyzing conflicts interactive and iterative. Moreover, to cope with interdependent relations, we adapt Satisfiability Modulo Theories (SMT) [27] techniques which can automatically determine the satisfiability of formulas expressed in first-order logic.

After conflict analysis, PIE is able to derive all the feasible solutions over finite domains, while the solutions satisfy user requirements and comply with enterprise policies. In PIE the service selection is modeled as a constraint satisfaction program (CSP) [28], and PIE automates the solving procedure using constraint programming (CP) techniques [29]. Interdependent relations and dependent domain constraints are encoded into a mathematical matrix which represents a selection problem; the matrix is then processed by a solver developed upon Choco [30], a Java CP library.

The technical details about PIE has been published as a technical report [23].

D. Cloud Intelligence

PIE generates a set of cloud service selections, and subsequently the service selections are fed to the Cloud Intelligence module as shown in Figure 1. Cloud Intelligence (CI) module is the key component in this framework to assess the solutions from PIE on the basis of multiple criteria including the quantitative metrics for cloud services, collective intelligence from intranet and social media, internal users' rating and feedback, and individual user's preferences.

CI takes several criteria into consideration to work out the optimal cloud service recommendation and selection for the users. Several MCDM algorithms can be applied to CI module to maximize the user's utility function based on the business requirements by setting the weights for various criteria from multiple sources connected to CI module in Figure 1. The quantitative metrics for cloud services including price, SLA, response time, etc., have been discussed in Section III-B as one source of criteria. The quantitative metrics provide objective indicators to measure the cloud services. Another source of criteria for cloud service recommendation and selection is social media which has been increasingly popular nowadays. Social media's power has been shown when users express themselves on social media in a timely and democratized way. In this framework the collective intelligence on cloud services are generated from social media and corporate intranet by mining and discovering the sentiments on the cloud services. We will leverage on the sentiment analysis engine [31] and event detection techniques [32] to improve the accuracy of cloud service recommendation and selection in CI.

A rating/feedback mechanism is put in place in the framework to allow the users sharing their experiences on the usage of various cloud services in the catalog. Such rating and feedback is considered coming from internal users of the platform, and therefore it becomes valuable assessment criterion in CI module. The user preferences on the cloud services and user patterns are another sources for CI to evaluate the cloud services.

In this framework, the solutions generated from PIE are evaluated on multiple criteria provided from different sources as shown in Figure 1, we employ MCDM techniques on CI module to generate the cloud service recommendation and selection for the users. The enterprises have the flexibility to adjust the weights for various criteria in this framework according to their preference on selecting cloud services.

IV. Conclusion and Future Work

The high level design of a systematic framework for enterprises on cloud service recommendation and selection in a hybrid cloud management platform is proposed in this paper. The key component in the framework, Policy Implementation Engine, has been successfully prototyped based on formal verification and constraint solving techniques, and it provides automated conflict detection and explanations of detected conflicts to identify problematic user requirements. The Cloud Intelligence in the framework leverages the MCDM techniques to generate cloud service recommendation and selection based on multiple criteria, e.g., quantitative metrics for cloud services, user preferences, collective intelligence from intranet and social media, and internal rating/feedback system.

Future work consists of working out actual use cases to evaluate the effectiveness of the cloud service recommendation and selection using proposed approaches, and further investigation on the MCDM algorithms used in the framework.

References

[1] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility," *Future Generation Computer Systems*, vol. 25, no. 6, pp. 599–616, 2009.

[2] P. Banerjee, R. Friedrich, C. Bash, P. Goldsack, B. A. Huberman, J. Manley, C. D. Patel, P. Ranganathan, and A. Veitch, "Everything as a service: Powering the new information economy," *IEEE Computer*, vol. 44, no. 3, pp. 36–43, 2011.

[3] B. Rochwerger, D. Breitgand, A. Epstein, D. Hadas, I. Loy, K. Nagin, J. Tordsson, C. Ragusa, M. Villari, S. Clayman, E. Levy, A. Maraschini, P. Massonet, H. Muñoz, and G. Tofetti, "Reservoir - when one cloud is not enough," *IEEE Computer*, vol. 44, no. 3, pp. 44–51, 2011.

[4] D. S. Milojicic, I. M. Llorente, and R. S. Montero, "OpenNebula: A cloud management tool," *IEEE Internet Computing*, vol. 15, no. 2, pp. 11–14, 2011.

[5] S. Yan, B. S. Lee, G. Zhao, D. Ma, and P. Mohamed, "Infrastructure management of hybrid cloud for enterprise users," in *Proceedings of the 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management: Standards and the Cloud*, Paris, France, Oct. 2011.

[6] Apache Deltacloud. [Online]. Available: http://deltacloud.apache.org

[7] "Cloud Infrastructure Management Interface (CIMI) Primer," Version 1.0.0d, Distributed Management Task Force (DMTF), 2012. [Online]. Available: http://dmtf.org/sites/default/files/standards/documents/DSP2027_1.0.0d.pdf

[8] AWS Marketplace. [Online]. Available: https://aws.amazon.com/marketplace

[9] Saleforce.com's appexchange. [Online]. Available: http://sites.force.com/appexchange

[10] VMware's solution exchange. [Online]. Available: http://www.vmware.com/appliances

[11] The Google Apps Marketplace. [Online]. Available: https://www.google.com/enterprise/marketplace

[12] HP Cloud Service Catalog. [Online]. Available: https://coral.saas.hp.com

[13] CloudSurfing. [Online]. Available: http://www.cloudsurfing.com

[14] DreamSimplicity. [Online]. Available: http://www.dreamsimplicity.com

[15] S. Yan, B. S. Lee, and S. Singhal, "A model-based proxy for unified iaas management," in *Proceedings of the 4th International DMTF Academic Alliance Workshop on Systems and Virtualization Management: Standards and the Cloud*, Ontario, Canada, Oct. 2010.

[16] P. Wright, T. J. Harmer, J. Hawkins, and Y. L. Sun, "A commodity-focused multi-cloud marketplace exemplar application," in *Proceedings of the 4th International Conference on Cloud Computing*. IEEE, 2011, pp. 590–597.

[17] S. K. Garg, S. Versteeg, and R. Buyya, "Smicloud: A framework for comparing and ranking cloud services," in *Proceedings of the 4th International Conference on Utility and Cloud Computing*. IEEE Computer Society, 2011, pp. 210–218.

[18] L. Wu, S. K. Garg, and R. Buyya, "SLA-based resource allocation for software as a service provider (SaaS) in cloud computing environments," in *Proceedings of the 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*. IEEE, 2011, pp. 195–204.

[19] Z. ur Rehman, F. K. Hussain, and O. K. Hussain, "Towards multi-criteria cloud service selection," in *Proceedings of the 5th International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing*, ser. IMIS'11. IEEE, 2011, pp. 44–48.

[20] B. Rochwerger, D. Breitgand, E. Levy, A. Galis, K. Nagin, I. M. Llorente, R. S. Montero, Y. Wolfsthal, E. Elmroth, J. A. Cáceres, M. Ben-Yehuda, W. Emmerich, and F. Galán, "The reservoir model and architecture for open federated cloud computing," *IBM Journal of Research and Development*, vol. 53, no. 4, p. 4, 2009.

[21] W. Zeng, Y. Zhao, and J. Zeng, "Cloud service and service selection algorithm research," in *Proceedings of the first ACM/SIGEVO Summit on Genetic and Evolutionary Computation*, ser. GEC '09. ACM, 2009, pp. 1045–1048.

[22] Cloudyn. [Online]. Available: http://www.cloudyn.com

[23] C. Chen, S. Yan, G. Zhao, B. S. Lee, and S. Singhal, "A systematic framework enabling automatic conflict detection and explanation in cloud service selection for enterprises," HP Laboratories, Tech. Rep. HPL-2012-104, May 2012. [Online]. Available: http://www.hpl.hp.com/techreports/2012/HPL-2012-104.pdf

[24] J. S. Dyer, P. C. Fishburn, R. E. Steuer, J. Wallenius, and S. Zionts, "Multiple criteria decision making, multiattribute utility theory: the next ten years," *Management Science*, vol. 38, no. 5, pp. 645–654, 1992.

[25] M. Ehrgott and X. Gandibleux, *Multiple Criteria Optimization: State of the Art Annotated Bibliographic Surveys*. Springer, 2002.

[26] D. Catteddu, "Security and resilience in governmental clouds," European Network and Inforamtion Security Agency (ENISA), Tech. Rep., 2011.

[27] C. Barrett, R. Sebastiani, S. A. Seshia, and C. Tinelli, "Satisfiability modulo theories," in *Handbook of Satisfiability*. IOS Press, 2009, ch. 26, pp. 825–885.

[28] S. Russell and P. Norvig, "Constraint satisfaction problems," in *Artificial Intelligence: A Modern Approach*, 2nd ed. Prentice Hall, 2002, ch. 5.

[29] K. Apt, *Principles of Constraint Programming*. Cambridge University Press, 2003.

[30] C. Team, "Choco: an open source java constraint programming library," Ecole des Mines de Nantes, Tech. Rep. 10-02-INFO, 2010.

[31] M. Castellanos, R. Ghosh, Y. Lu, L. Zhang, P. Ruiz, M. Dekhil, U. Dayal, and M. Hsu, "Livepulse: tapping social media for sentiments in real-time," in *WWW (Companion Volume)*, 2011, pp. 193–196.

[32] J. Weng, Y. Yao, E. Leonardi, and B. S. Lee, "Event detection in twitter," HP Laboratories, Tech. Rep. HPL-2011-98, July 2011. [Online]. Available: http://www.hpl.hp.com/techreports/2011/HPL-2011-98.pdf