Running Business Applications in the Cloud: A Use Case Perspective

Carmelo Ragusa and Antonio Puliafito

Faculty of Engineering, University of Messina,
Contrada di Dio, Messina, Italy
{cragusa,apuliafito}@unime.it

Abstract. Cloud computing leverages the use of abstracted resources. However, migrating an industrial application to well-known Cloud solutions such as EC2 might be complex and low level expertise is indeed needed. In this use case we present a methodology, based on practical experience matured on the ground, that allows service providers to enable complex applications to the RESERVOIR cloud infrastructure. We also show an example of how a complex business application, such as the SAP ERP 6.0, can be automatically fully deployed and scaled up and down as resource needs change, easing the use of a cloud system for service providers that might experience difficulties or have mental barriers to carry out such task.

Keywords: Cloud computing, Industrial applications, Use case, RESERVOIR, cloud enabling methodology.

1 Introduction

Today a growing number of companies, such as start-up and SMEs, is using Cloud computing to carry out their business, due to the easy requirements to be satisfied and to the fast high competitiveness gained at lower investment costs. However, Cloud computing is far from being definitive, since the always changing requirements are constantly modifying this paradigm. Enterprises are analysing the use of Cloud to carry out some of their processes, which add even more requirements. For example, many enterprise processes are time-critical, with secure and privacy requirements. Also the scale of enterprise applications is in the order of thousands of concurrent services, which is in contrast with SMEs that have a number of services orders of magnitude lower than enterprises. Another aspect in favour of SMEs for the Cloud adoption is that their system have a low number of functions, making no difference between data and logic, and therefore less complex to be "cloudified". This is not feasible for enterprises that have much more complex functions to be taken into account. The result is that Cloud solutions face difficulties to cope with both opposite class of customers and so the use of Cloud within enterprises is still low. However, the benefits in using Cloud are many and Cloud providers, pushed by the hype of the IT environment, are progressing fast to meet the new demands. Although there are

M.R. Guarracino et al. (Eds.): Euro-Par 2010 Workshops, LNCS 6586, pp. 595–602, 2011.
© Springer-Verlag Berlin Heidelberg 2011

efforts [1] to design new enterprise applications to be cloud oriented, they are far from being a reality. Thereby, adopting the Cloud for legacy applications can give enterprises advantages as well as speed up the usage of this paradigm.

The RESERVOIR solution provides an infrastructure able to run not only applications for SMEs, but also for large complex enterprises. The system offers capabilities such as rapid provisioning, elasticity, applications coexistence, federation and security. In this paper we focused on the first two features through a SAP ERP use case, showing how such complex application can be fully automatically deployed and scaled according to specific user requirements, without requiring any modification of the application. This will be an advantage not only for large enterprises using this application, but also for SMEs that are normally afraid of installing and configuring SAP systems due to their complexity. Finally, SAP itself can benefit because it can reach discouraged clients.

The reminder of the paper is organized as follows: section 2 discusses the related work. In section 3 we present the RESERVOIR solution along its main components. Next is discussed the use case, we give details about the application, the testbed and the scenarios performed. After, we describe the lesson learnt. Finally we draw the conclusions and discuss the future work.

2 Related Work

Different Cloud solutions (e.g. RightScale, Scalr, Flexiscale and Elastra) are available, offering high level functionalities on top of Cloud infrastructures such as Amazon EC2, Google App Engine and GoGrid. Such Cloud solutions try to fill the gap between Cloud infrastructures' offers and the Service Providers (SPs) that use them, by providing automation services to control the virtual resources assigned. However, their service definition is not comprehensive. In fact each Virtual Machine (VM) needs to be installed, configured and managed singularly, which restricts the service deployment. Moreover, the automation controls, such as auto-scaling, are too rigid due to the use of predefined monitored variables within the server templates offered, thereby not allowing the SPs to define their own service indicators. As result, SPs have to constantly monitor the VMs running state in order to instantiate or remove VMs to scale the service up or down.

In [2], the authors point out how elasticity is an important feature which allows SPs to save on costs of over-provisioning and risks of under-provisioning. Motahari et al [3] also highlights from a business perspective how such ability has a positive impact on the enterprise business processes management. In [4], the authors discuss the issues commented above and identify four goals, service abstraction level, automatic scalability, smart scaling and avoidance of Cloud vendor lock-in, thus reinforcing the importance to address those issues. On the other hand, migrating enterprises' IT systems to the Cloud, needs to be carefully studied. In [5] the authors investigated the implication to migrate an IT system to the Cloud. Few other works studied the impact of such migration from an enterprise perspective [6], showing also a real use case [7]. In the latter,

the authors show how migrating an IT infrastructure to a Cloud system can significantly reduce costs opposite to create and maintain an in-house solution.

The RESERVOIR project addresses the limits of current solutions, by providing a Cloud infrastructure where SPs can easily define their services that run in a fully automated fashion.

3 The RESERVOIR Solution

The RESERVOIR cloud solution [8][9][10] is shown in figure 1. SPs are external entities requiring resources for their applications based on high level business requirements.

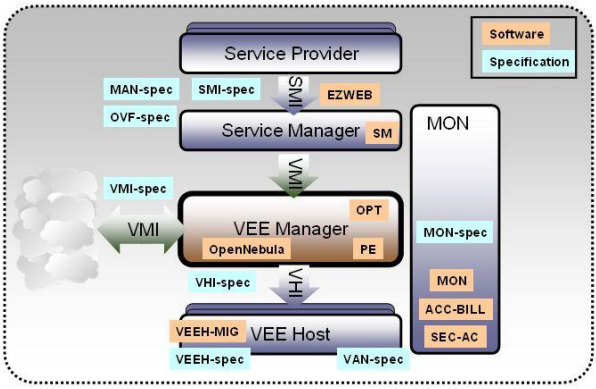

Fig. 1. RESERVOIR Cloud Architecture

A service in RESERVOIR is a set of Virtual Execution Environments (VEEs). Each VEE essentially wraps part of a SP's application. The SP specifies the terms of the service requested through a Service Manifest based on an extended version of the OVF [11].

Service Manager (SM). The SM is the module that interacting with the service providers, deals with the service deployment, checks elasticity rules and SLA compliance, and handles the service billing. The aim of this component is to automate a manual, complex and lengthy service delivery process. The SM processes the service requirements embedded within the manifest and determines the VEEs needed by the service, along their placement constrains based on cost, licensing, affinity, etc. SLAs are checked over the service life cycle and capacity is adjusted according to elasticity rules within the service manifest. This is accomplished by evaluating the application specific KPIs and deploying or removing the relative VEE instances. Finally, an accounting system within the SM processes the resources utilization of the service and creates bills according to post-paid or pre-paid billing models.

Virtual Execution Environment Manager (VEEM). The VEEM interacts with the SM and VEE Hosts (VEEHs) within the same site, and with remote VEEMs. The VEEM, based on the SM instructions, creates VEEs and places them in the VEEHs. Placement is done first to satisfy the SM constrains derived from the manifest, and second to optimize the site utilization according to the current local optimization policy such us of load balancing or power saving. Site policies can be easily plugged within the VEEM, so that each site can have its own optimization strategy. Cloud federation is also performed by this component. This feature allows to extend the cloud capacity over the physical limit of a single site. The VEEM can, within the manifest constraints, place VEEs across remote sites with which agreements have been created.

Virtual Execution Environment Host (VEEH). The VEEH interacts with the lower resources through the virtualization technology, such as XEN, KVM, etc, and also with the VEEM. The VEEH manages the VEEs abstracting the specific virtualization commands, into a common interface that the VEEM can use. In this way, the VEEM is unaware of the virtualization technology in place. VEEs can be deployed across different VEEHs and sites. This component also deals with network configurations. A specific module is dedicated to create Virtual Area Networks (VANs) for each application. This allows the application to communicate through a dedicated channel, independently of each VEE location. Therefore applications deployed in the same site are separated from each other.

Interfaces. The design of the RESERVOIR architecture facilitates the layers interoperability, by supporting open, generic and standard protocols and interfaces. Each layer can be implemented in a different way, but still able to intercat with each other. This approach will encourage new cloud enabling solutions. The Service Management Interface (SMI) allows service providers to access the RESERVOIR solution through the use of the OVF manifest. This means that different RESERVOIR providers use a common business requirements language. The VEE Management Interface (VMI) facilitates the use of different VEEMs within the RESERVOIR stack allowing to experiment with different management strategies. This interface allows also the cross-site communication needed to support the federation feature. Finally, the VEE Host Interface (VHI) supports the use of new virtualizaion technologies.

Monitoring Framework. The monitoring framework stretches across all layers. Its task is to pass information from lower to higher layers, in order for each component to take due action at runtime [12]. Applications running locally as well as across sites need to be monitored. Data can come from physical and virtual resources. For example, the SM needs to continuously check that elasticity rules and SLAs are satisfied. Therefore, probes have been developed and embedded within each VEE, to collect high level specific service data in the form of KPIs. Also, other probes interacting with the hypervisor and collecting CPU, memory and network usage were developed.

4 Use Case

The SAP Business Application. The use case presented in this paper is based on the SAP ERP 6.0, which is part of the SAP Business Suite. The ERP facilitates the flow of information among all business functions within a company. The suite runs over SAP NetWeaver Application Server (NWAS), a three tier architecture shown in figure 2. The application layer is made of two elements:

- **Central Instance (CI):** only one CI is present in a system and implements many services such as dialog, update, batch, spool, gateway, message server and enqueue server. Also, it provides higher level locking mechanism on the message server. It has high availability requirements and cannot be scaled out.
- **Dialog Instance (DI):** it processes user requests or Remote Function Call (RFC) from remote systems. Requests are handled in the form of work processes which can be dynamically parallelized. It can be scaled out.

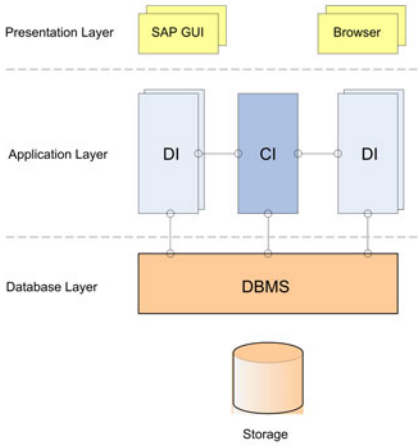

Fig. 2. SAP NETWeaver 3 tier architecture

The SAP NWAS can be configured either as 2-tier or 3-tier system. In a 2-tier system Database and Application layers form a single layer. A system deals with growing load by dynamically scaling up the DI instances. In our case, the NWAS was configured in a 2-tier fashion, where the CI and database were installed within the same VEE while the DI in a different VEE. As Presentation layer a proxy was installed within a VEE. The CI VEE requested 4 CPUs and 7GB of RAM, and its image was 120GB growing to 150GB with the swap area. The DI VEE requested 2 CPUs and 3GB of RAM, with an image of 1.5GB growing to 10GB with the swap area. The proxy requested 2 CPUs and 2GB of RAM. These VEEs requirements were coded within the service manifest.

Testbed Configuration. The RESERVOIR testbed is made of 4 sites, geographically distributed at UniMe (University of Messina, Italy), Umea (Umea University, Sweden), Thales (France) and IBM (Israel). For the presented scenario, the Umea site was used since the SAP application fitted in it. The Umea site is composed of 2 VEEHs and a management node running the SM and VEEM components. Each physical machine is equipped with Quad Xeon X3330 (2.66GHz), 8G RAM, and 160GB of Disk. A storage node stores the images.

4.1 Scenarios

Rapid Provisioning. The goal of this scenario is to demonstrate an automated full deployment of the SAP ERP 6.0 system. The system starts from an OVF descriptor, containing the details of the SP service, such as location of the images, minimum number of instances, VEE description (hypervisor, number of CPUs, memory, network) and elasticity rules to scale up and down the system. The SM contains an OVF parser [11] which extracts all necessary information to create, for each VEE, a descriptor file needed by the VEEM for the deployment, and a configuration file that will be used by the VEE itself. The configuration file, created as an ISO image, is mounted on the CD drive of the VEE that is then mounted by the VEEH. Next, the configuration file is used by the Activation Engine, within the image itself, that configures the VEE after the boot phase. Tiers need to know where to contact others tiers of the application, and the activation engine is used for this purpose. Figure 3 shows the process so far described.

As result, once deployed the CI (3-5 minutes from boot to running), the DI (3 minutes) and the SAPProxy (1 minute) get connected through a dedicated VAN, and the SAP application is fully configured. The result is a complex multi-tier business application, fully deployed within the RESERVOIR infrastructure, without any human intervention.

Elasticity. After the system has been deployed, 10000 concurrent users were simulated in order to generate a load that the system could not manage.

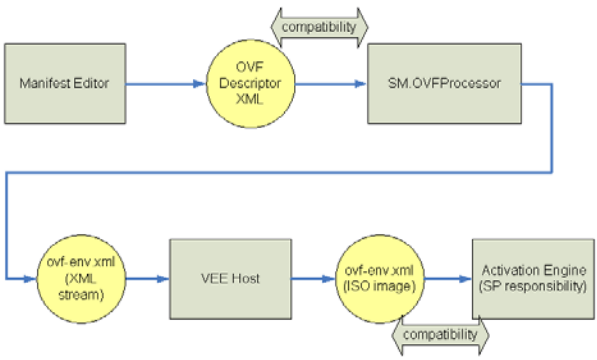

Fig. 3. SAP manifest processing

This was done through a load generator within the SAPProxy. The number of users represents the KPI for the SAP application, which was reported up to the SM, through the probe within the SAPProxy image that sends this data to the monitoring system. Since the Manifest contained elasticity rules to scale the system in such situations, a further DI was created and added to the SAP system, so that all users could be served. After a while the number of concurrent users was reduced by the load generator, which led the extra allocated DI to be removed by the SM.

5 Lesson Learnt

The SAP use case represents a complex class of applications, an enterprise multi-tier system, which allowed us to develop a methodology, based on our practical experience, to enable applications within the RESERVOIR cloud:

- **Virtualization of the application:**
  - embed the application in one or more VMs;
  - verify in-house that the application works;
- **Development of the OVF manifest:**
  - map the virtualized application into the manifest;
  - identify the valuable KPIs for the application;
  - create the appropriate elasticity rules;
- **Development of the activation engine:** configure each VM of the application appropriately;
- **Development of application's probes:** create the appropriate probes to push out the right KPIs from the VEEs;
- **Testing:** tests on deployment and elasticity in a real RESERVOIR environment.

6 Conclusions

This paper presented a complex business use case, the SAP ERP 6.0, running over the RESERVOIR cloud infrastructure. Two different scenarios were presented, the rapid provisioning and the elasticity. The former showed how the SAP ERP could be automatically deployed, starting from an OVF-based Manifest. The role of the Activation Engine to internally configure the VM, was also described. The latter scenario showed how the SAP application could be automatically scaled up and down, accordingly to the KPI. The process to bring this use case on the RESERVOIR cloud system allowed us to develop a methodology, which we formalized in a set of practical steps, that each SP can follow to enable its application to use the RESERVOIR solution.

Next we are going to demonstrate more features of the RESERVOIR system: *applications coexistence*: deploying more applications within the same site, in a transparent and secure fashion; *federation*: extending the infrastructure by setting agreements with other sites; and *cross sites live-migration* of VMs.

Acknowledgments. The research leading to the results presented in this paper has received funding from the European Unions seventh framework programme (FP7 2007-2013) Project RESERVOIR under grant agreement number 215605.

References

1. Chappell, D.: A short introduction to cloud platforms, an enterprise-oriented view. White paper, Chappell & Associates, 13 pages (August 2008)
2. Armbrust, M., Fox, A., Griffith, R., Joseph, A.D., Katz, R.H., Konwinski, A., Lee, G., Patterson, D.A., Rabkin, A., Stoica, I., Zaharia, M.: Above the clouds: A berkeley view of cloud computing. Tech. Rep. UCB/EECS-2009-28, EECS Department, University of California, Berkeley (February 2009)
3. Motahari-Nezhad, H.R., Stephenson, B., Singhal, S.: Outsourcing business to cloud computing services: Opportunities and challenges. Tech. Rep. HPL-2009-23, HP Laboratories (February 2009)
4. Rodero-Merino, L., Vaquero, L.M., Gil, V., Galn, F., Fontn, J., Montero, R.S., Llorente, I.M.: From infrastructure delivery to service management in clouds. Future Generation Computer Systems, Corrected Proof (2010) (in press)
5. Sriram, I., Khajeh-Hosseini, A.: Research agenda in cloud technologies. CoRR, abs/1001.3259 (2010)
6. Khajeh-Hosseini, A., Sommerville, I., Sriram, I.: Research challenges for enterprise cloud computing. CoRR, abs/1001.3257 (2010)
7. Khajeh-Hosseini, A., Greenwood, D., Sommerville, I.: Cloud migration: A case study of migrating an enterprise it system to iaas. CoRR, abs/1002.3492 (2010)
8. Rochwerger, B., Breitgand, D., Levy, E., Galis, A., Nagin, K., Llorente, I.M., Montero, R., Wolfsthal, Y., Elmroth, E., Caceres, J., Ben-Yehuda, M., Emmerich, W., Galan, F.: The reservoir model and architecture for open federated cloud computing. IBM Journal of Research and Development 53 (2009)
9. Rochwerger, B., Galis, A., Levy, E., Cáceres, J.A., Breitgand, D., Wolfsthal, Y., Llorente, I.M., Wusthoff, M., Montero, R.S., Elmroth, E.: Reservoir: management technologies and requirements for next generation service oriented infrastructures. In: IM 2009: Symposium on Integrated Network Management, Piscataway, NJ, USA, pp. 307–310. IEEE Press, Los Alamitos (2009)
10. Rochwerger, B., Galis, A., Breitgand, D., Levy, E., Cáceres, J.A., Llorente, I.M., Wolfsthal, Y., Wusthoff, M., Clayman, S., Chapman, C., Emmerich, W., Elmroth, E., Montero, R.S.: Design for future internet service infrastructures. In: Towards the Future Internet - A European Research Perspective, pp. 227–237. IOS, Amsterdam (2009)
11. Galán, F., Sampaio, A., Rodero-Merino, L., Loy, I., Gil, V., Vaquero, L.M.: Service specification in cloud environments based on extensions to open standards. In: COMSWARE 2009: Communication System Software and Middleware, pp. 1–12. ACM, New York (2009)
12. Clayman, S., Galis, A., Toffetti, G., Rodero-Merino, L., Vaquero, L.M., Nagin, K., Rochwerger, B.: Monitoring service clouds in the future internet. In: Towards the Future Internet - Emerging Trends from European Research, pp. 1–12. IOS, Amsterdam (2010)