

Analysing Scalability Strategies for Service Choreographies on Cloud Environments

Raphael Gomes^{1,2(✉)}, Fabio Costa¹, and Ricardo Rocha¹

¹ Instituto de Informática, Universidade Federal de Goiás, Goiânia, Brazil
`{fmc,ricardo}@inf.ufg.br`

² Instituto Federal de Goiás, Itumbiara, Brazil
`raphael.gomes@ifg.edu.br`

Abstract. Scalability is one of the major advantages brought by cloud computing environments. This advantage can be even more evident when considering the composition of services through choreographies. However, when dealing with applications that have quality of service concerns scalability needs to be performed in an efficient way considering both horizontal scaling - adding new virtual machines with additional resources, and vertical scaling - adding/removing resources from existing virtual machines. By efficiency we mean that non-functional properties must be offered in the choreographies while is made effective/improved resource usage. This paper discusses scalability strategies to enact service choreographies using cloud resources. We present efforts at the state of the art technology and an analysis of the outcomes in adopting different strategies of resource scaling. We also present experiments using a modified version of CloudSim to demonstrate the effectiveness of these strategies in terms of resource usage and the non-functional properties of choreographies.

Keywords: Cloud computing · Scalability · Choreography · Auto-scaling

1 Introduction

The provision of quality of service (QoS) is one of the main challenges in cloud computing [5], since this paradigm must provide assurances that go beyond the typical maintenance activities and must provide high reliability, scalability and autonomous behaviour. Many QoS aspects of an application are related with the scalability provided by the hardware resources used to deploy it. As a matter of fact, cloud environments are increasingly used due to their elasticity and the illusion of infinite resources. Increasing degrees of scalability are achieved through the automated management of resources, typically using horizontal scaling, which means changing the amount of resources used, by adding or removing virtual machines (VM) according to policies related to the use of such resources or non-functional properties of the application.

Another strategy for scalability is the use of vertical scaling, i.e. on-the-fly changing of the amount of resources allocated to an already running VM instance, for example, allocating more physical CPU time to a running virtual machine. In a complementary manner, we can have a hybrid approach where we increase both the number and the configuration of virtual machines.

Although there are different scalability strategies they must be used in an efficient way, regarding the consumption of resources. This is due to the fact that a poor management of resources can result in unnecessary spending in the case of public clouds, as well as problems related to energy consumption and loss of investment in the case of private clouds. Another issue that must be taken into account is the QoS offered to applications since some functionalities may not be useful if certain non-functional attributes are not guaranteed [8].

These challenges are even more evident in the so called Future Internet, which results from the evolution of the current Internet, in combination with the Internet of Content [9], the Internet of Services [24] and the Internet of Things [2]. In this new paradigm there are thousands of services belonging to different organizations that have to cooperate with each other in a distributed and large scale environment. This integrated view of services highlighted some problems that were not readily apparent in previous integration efforts, which hardly reached the scale that systems of web services now have [30].

Keeping centralized coordinators for these new types of applications is infeasible due to requirements like fault tolerance, availability, heterogeneity and adaptability. For this reason, the most promising solution may be the organization of decentralized and distributed services through choreographies. Choreographies are service compositions that implement distributed business processes, typically between organizations in order to reduce the number of exchanged messages and distribute business logic, without the need for centralized coordinators, since each service “knows” when to perform its operations and which other services it must interact with [3].

This paper discusses the state of the art in providing scalability for cloud-based service choreographies, considering both technologies and cloud providers. We discuss the outcomes of using cloud environments and the main advantages and disadvantages of adopting different scalability strategies to enact choreographies on cloud resources. We strengthen this discussion with some preliminary evaluation of these strategies. Although this article does not address aspects related to the implementation of choreographies, the analysis presented here can be used as input to different approaches regarding choreography execution in cloud, as well as general applications. The remaining of this paper is as follows: Sect. 2 presents a motivating example; Sect. 3 discusses how actual virtualization technologies and cloud providers handle scalability strategies, while Sect. 4 presents some results of the evaluation of these strategies in choreography enactment. Finally, Sect. 5 discusses related work and Sect. 6 presents the final remarks.

2 Motivating Example

Media sharing is one of the main Internet applications [21]. This type of application was driven by the increasing use of social networks and content sharing

platforms such as YouTube, Instagram and Facebook, and its growth brings scalability problems, with increasing demands for data storage and transfer, and the pressure to deliver faster service and other quality attributes. Cloud computing is therefore an increasingly used alternative for resource providers to circumvent these problems. Therefore, in this section we will explore some scenarios to illustrate the complexities involved in the management of scalability issues.

Let us suppose a fictitious organization that uses a public cloud provider or a datacenter (using some virtualization technology) to obtain resources for its applications. One of the applications consists in a choreography of services for media sharing on the web. This application comprises a service for media upload that communicates with other two services: one to perform media storage and another to perform media indexing. There is also a service that provides the website front end.

The services of the application have associated quality of service requirements: initially the upload service must handle at least 100 concurrent requests; data storage must be performed in a secure environment; and indexing overhead should be less than 1 s for at least 90% of the requests. Based on these requirements and on the expected demand, let us suppose that in order to enact the choreography, it is necessary to allocate one VM for the upload service, as well as using a scalable architecture for media storage, a relational database for data indexing, and another VM for the front end. The services and resources of the initial scenario are illustrated in Fig. 1(a).

After choreography deployment and application execution starts, suppose there is a considerable increase in demand. This behaviour is common in many web applications, which typically starts with only a handful of users but quickly grow to reach thousands and even millions of users. As an example, Facebook has an average growth of 250,000 new users per day [12].

In addition to increased demand, in our scenario another requirement was raised - the viewing of media in various formats. Accordingly, it is necessary to convert the original media, which is a intensive processing task. Thus, two new services must be added to the choreography: one to perform media processing before storage and another to control the queue for this. Furthermore, aiming to increase competitiveness, the upload QoS was modified, aiming to be able to handle ten times more requests concurrently. To meet this new scenario, it is necessary to review the initial resource allocation.

In the additional resource allocation we can adopt the strategies cited before: the first one is to do horizontal scaling. Accordingly, we can create other VM instances and get something like allocation *A* in Fig. 1(b). On the other hand, we can use vertical scaling and keep the number of resources but increasing their configuration, as in the allocation *B* in Fig. 1(b).

The main problem in this scenario is to decide which scalability strategy is the preferred option given the quality of service requirements and the cloud provider or technology features, e.g. performance, security, cost, etc. For instance, at a first glance horizontal scaling is a good choice for the media processing service due concurrency issues but what are the outcomes of adopting this strategy instead




Fig. 1. Choreography deployment for the media sharing application.

of vertical scaling? In addition we can even use both strategies by allocating more instances with an increased configuration. Another point is ability to determine the overhead as well as the ability to adopt each of these strategies when using a the given provider/technology.

3 Scalability Strategies in Cloud Platforms

Current virtualization technologies allow the addition of new VM instances as well as the re-dimensioning of running VMs. For horizontal scaling, experiments have shown different values for VM startup time [4, 17], although on average it takes about 1 min, while vertical scaling allows to double the processing power in less than 1 s [34]. Gong et al. [15] indicate that changes in the amount of CPU take on average 120 ± 0.55 ms.

Another positive factor for the adoption of a hybrid approach rather than the commonly adopted practice of only scaling resources horizontally is that vertical scaling is quite advantageous in some cases, as Dawoud et al. [10] demonstrate. According to these authors, a web server running on a vertically scaled VM offers better performance than a web server running multiple VMs, i.e. a VM with 4 cores implies a lower response time compared to 4 VM instances running in parallel.

Table 1 presents details on automated resource scaling for some of the main providers and cloud technologies. As can be seen, the majority of them do not support automatic scaling but provide APIs that allow one to perform this task.

Table 1. Scalability in cloud providers/technologies.

Cloud Provider/Technology	Automatic horizontal scaling	Automatic vertical scaling
Amazon (aws.amazon.com)	yes	no
Windows Azure (www.windowsazure.com)	yes	no
Google App Engine (developers.google.com/appengine)	yes	no
Google Compute Engine (cloud.google.com/products/compute-engine)	yes	no
Rackspace (www.rackspace.com)	no	yes
Flexiscale (flexiscaletechnologies.co.uk)	no	yes
GoGrid (www.gogrid.com)	no	no
Joyent Cloud (www.joyent.com)	no	yes
Eucalyptus (www.eucalyptus.com)	no	no
Xen (www.xenproject.org)	no	yes

By using the *Auto Scaling* and *Cloud Watch* in Amazon EC2, it is possible to define policies for VM creation and destruction. The average startup time for a new instance on this provider is between 2 and 10 min, though this time is close to 100s for instances running Linux [20]. Similarly, VM startup time in Windows Azure is around 10 min, although different requests for VM creation may take up different amounts of time. Experiments in [17] indicate a delay of 4 min between the startup time of the first and fourth instance using Azure. Some cloud providers like Google App and Compute Engine¹ offer horizontal auto scaling, although it is not possible to manage its operation.

Regarding vertical scaling, in Rackspace, Flexiscale and Joyent Cloud it is possible to do scaling of processor but this requires VM reboot. CPU scaling can be performed until it reaches the full capacity of the underlying physical machine [32]. In Xen we can also have scaling of memory, which is performed by a process known as *memory ballooning*, which allows changing of the amount of memory to a VM while it is running. A similar mechanism is offered for CPU scaling [27].

¹ In the Google Compute Engine horizontal auto scaling is implemented as an application from App Engine.

In addition to the APIs offered by cloud providers is possible to use frameworks to achieve automated scaling of resources. Table 2 [13] shows some of these frameworks, together with the providers they support and details on the scaling capacity.

Table 2. Frameworks for cloud resource management.

Framework	Supported providers	Scaling strategy
Cloudify (www.cloudifysource.org)	Amazon, OpenStack, Azure, HP Cloud, Rackspace	Automated scaling based on metrics related to VM configuration and number of instances
Cloud Foundry (www.cloudfoundry.com)	Amazon, OpenStack, Rackspace, Eucalyptus	Manual change in the number of VMs associated with an application
Scalr (www.scalr.com)	Amazon, OpenStack, Rack- space, Nimbula, Eucalyptus, IDC Frontier, CloudStack, Cloud Foundry	Automated scaling of infrastructure and database when they are overloaded or when scheduling is done
OnApp (onapp.com)	Xen, KVM, VMware and other smaller public provider	Vertical scaling with automated VM migration when the physical machine does not have enough resources

Besides scaling, it is necessary to take into account the time required to perform VM migration from one physical machine to another when the existing resources are insufficient. Considering Xen for instance, VM migration requires the transfer of all memory. However, the migration mechanism hides the latency by continuing running the application on the original VM while the memory contents are transferred. Experiments in [25] show that the migration of a VM with 800 MB of memory on a LAN using Xen caused unavailability of 165 ms to 210 ms and increased application execution time in 17–25 s. However, during this period throughput decreases only 12 %. The results in [31] show that, in an instance of an almost overloaded system (serving 600 concurrent users), migration causes significant downtime (about 3 s) but the 99th SLA percentile could be met, i.e. 99 % of more critical SLA could be met.

Based on the above, although there are no records of accurate results for all technologies in the literature, we can state that the achievement of scaling in two dimensions (vertical and horizontal) is feasible. Having said that, we now present some experiments to evaluate how these strategies work in the context of service choreographies.

4 Evaluating Scalability Strategies to Service Choreographies

We performed some preliminary experiments to analyse the scalability strategies and evaluate how it interferes in the resource usage and the QoS offered to choreographies. Through this experiment, we expect to answer the following research questions:

- Does the cost associated with the allocation of resources justify potential advantages obtained?
- Which one is the best strategy for resource scaling in a choreography enactment?
- How different scaling policies interfere in the usage of resources and non-functional characteristics of services choreographies?

The experiments were performed using simulations through a modified version of CloudSim [7]. Initially we present the modifications implemented and then we describe the experiments and their results. To design and describe the experiments we follow our structure proposed by Wohlin et al. [33].

4.1 Simulator

CloudSim [7] is an extensible toolkit that enables modelling and simulation of clouds, as well simulation of policies for resource provisioning. This simulator is usually used to investigate the infrastructure design decisions by analysing different configurations [6]. Cloud providers are modelled in the simulator as *Datacenters* receiving service requests. These requests are elements encapsulated in *VMs* that need to allocate shared processing power in a given *Host* in the datacenter. The *VMScheduler* component is responsible for scheduling the host that runs each VM.

Applications running in the cloud environment are represented as a set of *Cloudlets*, which store execution data as request size in millions of instructions (MI); and utilization modes of CPU, memory and bandwidth. The *Datacenter-Broker* component is responsible for simulation and management of cloudlets, and to configure their policies of resource management.

CloudSim has some limitations that hinder the simulation of our scenario. For example, it does not enable the simulation of automatic resource provisioning and it simulates just a single service. Thus, we implemented an extension of CloudSim to overcome these limitations and enable the simulation of choreography enactment. The implementation of auto scaling mechanism was made using a simplified version of the model proposed by Amazon for its services *CloudWatch* and *Auto Scaling*. According to this model, it is possible to establish metrics for resource monitoring, and policies for VM manipulation.

For representation of choreographies, we adopted the semantics of Web Service Choreography Description Language (WS-CDL) [19], an XML-based language to describe collaboration between multiple stakeholders in a process of

business. WS-CDL is a W3C recommendation detailed at <http://www.w3.org/TR/2005/CR-ws-cdl-10-20051109>.

The modifications made in CloudSim not influence the core operation, since we added only metrics collection and resource allocation policies. Thus, the results remain equivalent to a real scenario.

4.2 Simulation of Resource Allocation Policies Applied to a Service Choreography

We performed simulations of running the media sharing application in a cloud environment adopting different policies for resource allocation. In those simulations, we aimed to analyse how those policies impact in resource usage and the application QoS.

Experiment Design. Our experiment evaluates how different resource scaling policies interfere in the usage of cloud resources and the resulting choreography's QoS.

Experiment Planning. We evaluate the scalability by simulating horizontal, vertical and hybrid scaling policies. The simulated scenario is the enactment of the media sharing choreography (Sect. 2) in a single cloud provider. The following variables were analysed:

- **Latency:** mean enactment time of the choreography.
- **Usage (VM):** mean and variance of the load in the virtual machines.
- **Usage (Host):** mean and variance of the load in the hosts (physical machines).
- **AWRT:** Average Weighted Response Time, as proposed by Grimme et al. [16], which measures how much, on average, users should expect to have their requests met.
- **Execution Overhead:** how much, on average, the execution time differs from an estimated optimal value. This value is taken assuming that the request in question is the only one in the cloud, i.e., no other concurrent requests share the same resource.

The size of the media used as input in the simulation was obtained randomly, whereas values between 0 and 10 GB. The activities duration, expressed in MI, that composes the choreography was estimated according to the size of the media, considering the average connection speed in Brazil and values obtained with benchmarks [22, 23, 26].

We evaluated scenarios with different requests per second rates: 1, 5, 10 and 30. For the horizontal scaling we performed scenarios with 1, 5, 10 and 30 running VMs. We used VM configurations equivalent to the types of Amazon EC2 *m1.small*, *m1.medium*, *m1.large* and *m1.xlarge* for the vertical scaling. In the hybrid approach, both the amount and configuration of VMs have changed. For each simulation, we consider a fixed amount of resources allocated. The simulated cloud consists of 10 hosts with same configuration: equivalent to a machine Intel[®] Xeon[®] Processor X5570, 8 GB of RAM.

Experiment Execution. This experiment is based on the modified version CloudSim described above. The choreographies submitted in the simulation were generated and stored so that the same input sets were used in the three scaling strategies.

There was no need of any treatment regarding the validity of the data.

Analysis and Interpretation of Results. Each experiment starts with the same scenario of one VM type *small* and load of one request/s. To evidenciate the difference among results, charts were plotted on a logarithmic scale, except for the chart of strategy costs (Fig. 3).

The first variable taken was the latency. As can be seen in Fig. 2(a), when there is 1 req/s the increase in VM number only brings high gains when going from 1 to 5 VMs, with a decrease of 52.8% on average latency. With scaling to 10 VMs the gain decreases to 4.14% and after that no more gains are obtained. This is due the fact that there are few requests, since a similar behaviour is observed in horizontal scaling when there are 5 req/s. Therefore, the horizontal scaling can only be justified for large-scale scenes, such as 30 req/s, where the gain is always greater than 50% (Fig. 2(a)).

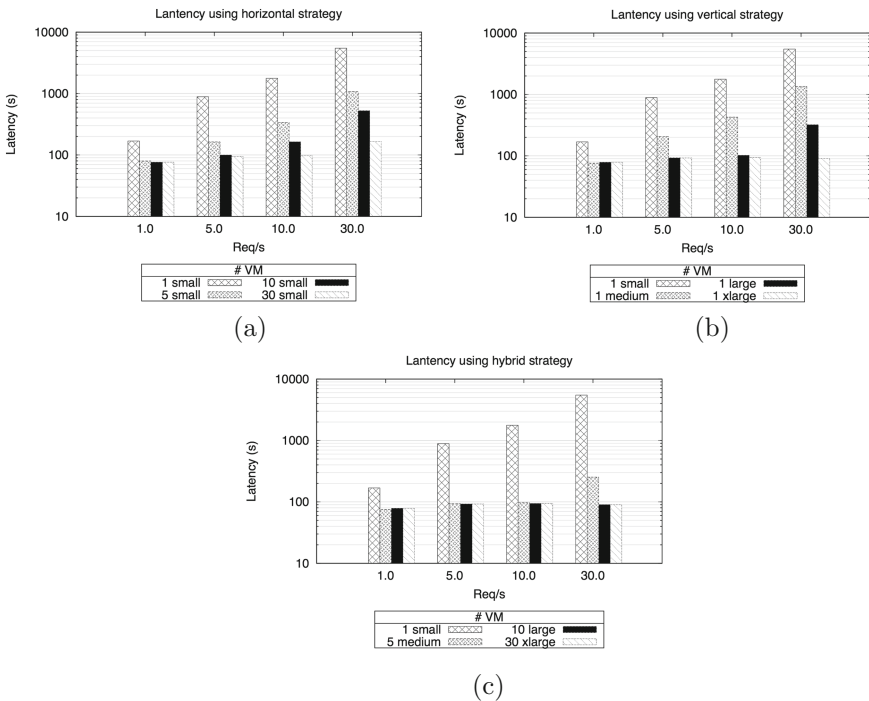


Fig. 2. Latency. (a) Horizontal strategy. (b) Vertical strategy. (c) Hybrid strategy.




Fig. 3. Cost of scaling. (a) Horizontal strategy. (b) Vertical strategy. (c) Hybrid strategy. (d) Comparative among strategies.

Comparing the horizontal and vertical approaches (Figs. 2(a) and (b)), we can see that, regarding latency, to use only 1 VM type *medium* is almost as satisfying as using 5 VMs type *small*. For types *large* and *xlarge* the vertical strategy is more advantageous than the horizontal strategy. This shows that the allocation of a larger number of VMs is not the best approach in this scenario. This is mainly because of the cost: considering a public cloud, the most extreme case (30 req/s) would require the expenditure of \$2.40 per hour in a horizontal approach and \$0.64 using the vertical strategy².

The hybrid strategy (Fig. 2(c)) shows that increasing the amount and configuration of VMs only bring benefit in the first modification, i.e. move from one VM type *small* to 5 VMs type *medium*. The only exception is the case of 30 req/s, for which the second modification also brought gains. Comparing the vertical and hybrid approaches (Figs. 2(a) and (b)) we found that the hybrid strategy has not brought gains in some cases. This reinforces the argument that changing the configuration of VMs can be a better strategy than allocate a greater amount of VMs.

² This estimation is using Amazon EC2 resources in São Paulo (Brazil) availability zone in February/2014.

To compare each scaling strategy, we defined the cost of each scaling, called *normalized cost*, by the expression $cost_per_hour / (QoS \times load)$, where *cost_per_hour* is the cost to implement the strategy, considering the number and type of VMs, *QoS* is the inverse of latency and *load* is the number of requests per seconds. Thus, a scaling strategy is better when the normalized cost decreases as the scale increases, for a same load. To a fair comparison among different strategies, we converted each strategy scale to a relative scale cost (per hour), considering that one small VM has a cost of one unit.

Figure 3 shows graphs of the normalized cost for each scaling strategy. Figure 3(d) shows a comparison between the cost of each strategy, on a load of 30 request/s. Figure 3(d) endorses the idea that the vertical strategy is the best strategy for the simulated scenario. Also, each strategy presents a point where it does not produce any benefit. As shown in Figs. 3(a) and (c), scaling in horizontal and hybrid approaches is only a beneficial approach at higher loads (close to 30 request/s).

The execution overhead and AWRT presented the same behaviour of latency, so the corresponding charts will not be shown or discussed. We also analyse the impact of scaling approaches on the use of resources, since the use of resources must be maximized in private clouds. Initially we analyse the use of physical machines. Since this factor depends solely on the amount and configuration of VMs created and there is no change in these attributes at runtime, the results are similar for all the requests rate and, so we put them together to facilitate comparison.

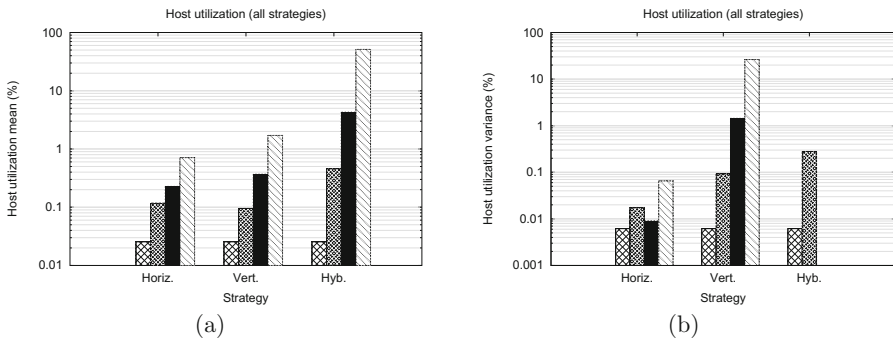


Fig. 4. Host usage. (a) Average. (b) Variance.

In Fig. 4 we can see the average and variance of utilization of physical machines. As showed in Fig. 4(a), the average host usage has no major changes when comparing the horizontal and vertical approaches, despite a much larger amount of VM created in the horizontal approach. Regarding the hybrid approach, only there was considerable use of resources in the latter case (30 VMs type *xlarge*), where the average utilization of physical machines was above 50%. The low utilization rates for the other cases is due to the large number of available resources in the cloud.

The usage variance (Fig. 4(b)) shows that resources were allocated almost uniformly. The only exception occurred in the vertical approach using VM type *xlarge* for which variance was 26.31 %.

Regarding the VM usage, we note that even with a greater number of instances, the horizontal approach has a higher average utilization (Fig. 5(a)). This was expected since, by having a more limited setting, VMs need to conduct a more intensive processing to complete the activities. On the contrary, the creation of 30 VMs makes the average utilization of VMs lesser than 2 % for 1, 5 and 10 req/s. Even at 30 req/s the usage is no more than 25 %, which shows that most of these resources are idle, representing a loss. On the other hand, use only 1 or 5 VMs (with same configuration) causes the average utilization be near 100 % when there are many requests.

The vertical approach also seemed to be the best alternative with respect to the use of resources. Mean usage (Fig. 5(b)) with the type *medium* was around 50 % in almost all cases. For the type *large* this value was lower, not exceeding 25 %. The use of type *xlarge* makes the average utilization be very low, less than 4 % in all cases. For the hybrid case (Fig. 5(c)) the scaling only improves substantially the usage from the first to the second simulated case (5 VMs type *medium*), since allocation of more resources decrease the average usage to a low percentage, no more than 1 %.

In the analysis of usage variance it makes no sense to take into account the cases with only one VM running. As a result, according to Figs. 5(d) and (e) we can see that there are major differences only in extreme cases where the amount of VMs is greater than or equal to 10. This happens because in the other cases there is a limited number of VM, which makes usage occurs almost uniformly.

With this experiment we concluded that the allocation of a greater number of VMs is not the best strategy when considering services choreographies. Our results suggest that, due to dependencies between choreography activities, the use of a strategy that minimizes the execution time may produce improved results. However it is still necessary a mechanism to decide which strategy is the best in each scenario, using adaptive algorithms that learn with the execution historic. For services choreographies we can even use different strategies to each services set. In a future experiment, we plan to evaluate how a choreography style and architecture may influence the gain of each scaling strategy.

5 Related Work

There have been significant research efforts on the analysis of virtualization technologies [18,31] or cloud providers performance [11,17], mainly related to effects of horizontal or vertical scaling on the perceived latency.

Vaquero et al. [29] survey various approaches for application scalability in clouds at three different levels: server, network and platform; they also discriminate the two types of scalability: horizontal and vertical.

Suleiman et al. [28] identify a series of research challenges related to the scalability of existing solutions. Their work aims to determine how much to scale,




Fig. 5. VM usage. Average on horizontal (a), vertical (b) and hybrid (c) strategy. Variance on horizontal (d) and hybrid (e) strategy.

taking into account automated scaling mechanisms and the costs associated with licensing, as well as the flexibility enabled by the size and type of resources that can be scaled. They also analysed how to scale, and which scaling strategy to choose, conducting a trade-off analysis between horizontal and vertical solutions.

Nevertheless none of these works considers scalability in clouds to enact service choreographies. As we pointed out above, there are some particularities that must be taken into account, like dependencies between services or concurrency issues. To the best of our knowledge there is no other work that considers this subject.

6 Final Remarks

In this paper we discuss how the main cloud providers and virtualization technologies provide scalability. By means of experiments carried out through simulation, we investigate the impacts of using three scalability strategies do enact service choreographies on cloud resources.

While vertical scalability is possible in principle for all applications, it largely depends on the service provider to offer the mechanisms to implement this type of 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 [1]. As pointed out the overhead of each strategy also needs to be taken into account. Horizontal scaling, for instance, requires on average 10 min to start a new VM instance, which may not be feasible in scenarios that involve real-time applications.

We evaluated some scalability scenarios to enact service choreographies. In our analysis we concluded that vertical scaling is the best option in terms of cost, resource usage and application QoS attributes for the application considered. This result, although not be general for various application areas, it has applicability in resource allocation in approaches that use cloud computing in the general case. For instance, a model-driven development process can manage the relation between the infrastructure and the actual application non-functional requirements choosing the most suitable scalability strategy to meet these requirements [14].

Nevertheless important scenarios with scalability patterns still need to be evaluated. In addition, a more precise analysis must be performed in order to provide elements that will enable a more effective scalability strategy. In particular, it is necessary to characterize the nature and behaviour of choreographies, including the characteristics of each individual service, using this information to refine the evaluation of scalability strategies. In our work the main focus was application load but this analysis can take into account other aspects such as service load and the use of public vs. private clouds.

References

1. Andrikopoulos, V., Binz, T., Leymann, F., Strauch, S.: How to adapt applications for the Cloud environment. *Computing* **95**(6), 493–535 (2013)
2. Atzori, L., Iera, A., Morabito, G.: The internet of things: a survey. *Comput. Netw.* **54**(15), 2787–2805 (2010)
3. Barker, A., Walton, C., Robertson, D.: Choreographing web services. *IEEE Trans. Serv. Comput.* **2**(2), 152–166 (2009)
4. Bellenger, D., Bertram, J., Budina, A., Koschel, A., Pfänder, B., Serowy, C., Astrova, I., Grivas, S., Schaaf, M.: Scaling in cloud environments. *Recent Researches in Computer Science* (2011)
5. Blair, G., Kon, F., Cirne, W., Milojicic, D., Ramakrishnan, R., Reed, D., Silva, D.: Perspectives on cloud computing: interviews with five leading scientists from the cloud community. *J. Internet Serv. Appl.* **2**(1), 3–9 (2011)

6. Caglar, F., An, K., Shekhar, S., Gokhale, A.: Model-driven performance estimation, deployment, and resource management for cloud-hosted services. In: Proceedings of the 2013 ACM Workshop on Domain-Specific Modeling, pp. 21–26. ACM (2013)
7. Calheiros, R., Ranjan, R., Beloglazov, A., De Rose, C., Buyya, R.: CloudSim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. *Softw. Pract. Exper.* **41**(1), 23–50 (2011)
8. Chung, L., do Prado Leite, J.C.S.: On non-functional requirements in software engineering. In: Borgida, A.T., Chaudhri, V.K., Giorgini, P., Yu, E.S. (eds.) *Conceptual Modeling: Foundations and Applications*. LNCS, vol. 5600, pp. 363–379. Springer, Heidelberg (2009)
9. Daras, P., Williams, D., Guerrero, C., Kegel, I., Laso, I., Bouwen, J., Meunier, J., Niebert, N., Zahariadis, T.: Why do we need a content-centric future Internet? Proposals towards content-centric Internet architectures. *Inf. Soc. Media J.* (2009)
10. Dawoud, W., Takouna, I., Meinel, C.: Elastic virtual machine for fine-grained cloud resource provisioning. In: Krishna, P.V., Babu, M.R., Ariwa, E. (eds.) *ObCom 2011, Part I. CCIS*, vol. 269, pp. 11–25. Springer, Heidelberg (2012)
11. Dejun, J., Pierre, G., Chi, C.-H.: EC2 performance analysis for resource provisioning of service-oriented applications. In: Dan, A., Gittler, F., Toumani, F. (eds.) *ICSOC/ServiceWave 2009*. LNCS, vol. 6275, pp. 197–207. Springer, Heidelberg (2010)
12. Facebook: Statistics (2013). <https://newsroom.fb.com>
13. Ferry, N., Rossini, A., Chauvel, F., Morin, B., Solberg, A.: Towards model-driven provisioning, deployment, monitoring, and adaptation of multi-cloud systems. In: *CLOUD 2013: IEEE 6th International Conference on Cloud Computing*, pp. 887–894 (2013)
14. Gomes, R., Costa, F., Bencomo, N.: On modeling and satisfaction of non-functional requirements using cloud computing. In: *Proceedings of the 2013 IEEE Latin America Conference on Cloud Computing and Communications* (2013)
15. Gong, Z., Gu, X., Wilkes, J.: *Press: predictive elastic resource scaling for cloud systems*. In: 2010 International Conference on Network and Service Management (CNSM), pp. 9–16. IEEE (2010)
16. Grimme, C., Lepping, J., Papaspyrou, A.: Prospects of collaboration between compute providers by means of job interchange. In: Frachtenberg, E., Schwiegelshohn, U. (eds.) *JSSPP 2007*. LNCS, vol. 4942, pp. 132–151. Springer, Heidelberg (2008)
17. Hill, Z., Li, J., Mao, M., Ruiz-Alvarez, A., Humphrey, M.: Early observations on the performance of Windows Azure. In: *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing*, pp. 367–376. ACM (2010)
18. Huang, X., Bai, X., Lee, R.M.: An empirical study of VMM overhead, configuration performance and scalability. In: *2013 IEEE 7th International Symposium on Service Oriented System Engineering (SOSE)*, pp. 359–366 (2013)
19. Kavantzias, N., Burdett, D., Ritzinger, G., Fletcher, T., Lafon, Y., Barreto, C.: *Web services choreography description language version 1.0*. W3C candidate recommendation, 9 (2005)
20. Li, A., Yang, X., Kandula, S., Zhang, M.: CloudCmp: comparing public cloud providers. In: *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, pp. 1–14. ACM (2010)
21. Miller, M.: *Cloud Computing: Web-Based Applications that Change the Way You Work and Collaborate Online*. Que Publishing, Indianapolis (2008)
22. Movavi: *Faster Performance with Cutting-Edge Tech* (2014). <http://www.movavi.com/videoconverter/performance.html>

23. MySQL: Estimating Query Performance (2014). <http://dev.mysql.com/doc/refman/5.0/en/estimating-performance.html>
24. Papadimitriou, D.: Future Internet - The cross-ETP vision document. European Technology Platform, Alcatel Lucent 8 (2009)
25. Ruth, P., Rhee, J., Xu, D., Goasguen, S., Kennell, R.: Autonomic live adaptation of virtual networked environments in a multidomain infrastructure. *J. Internet Serv. Appl.* **2**(2), 141–154 (2011)
26. Seagate: Savvio[®] 10K.5 SAS Product Manual (2012)
27. Senthil N.: Dynamic resource provisioning for virtual machine through vertical scaling and horizontal scaling. Ph.D. Dissertation. Department of Computer Science and Engineering, Indian Institute of Technology (2013)
28. Suleiman, B., Sakr, S., Jeffery, R., Liu, A.: On understanding the economics and elasticity challenges of deploying business applications on public cloud infrastructure. *J. Internet Serv. Appl.* **3**(2), 173–193 (2012)
29. Vaquero, L.M., Rodero-Merino, L., Buyya, R.: Dynamically scaling applications in the cloud. *SIGCOMM Comput. Commun. Rev.* **41**(1), 45–52 (2011)
30. Vincent, H., Issarny, V., Georgantas, N., Franceschini, E., Goldman, A., Kon, F.: CHOReOS: scaling choreographies for the internet of the future. In: *Middleware'10 Posters and Demos Track*, p. 8. ACM (2010)
31. Voorsluys, W., Broberg, J., Venugopal, S., Buyya, R.: Cost of virtual machine live migration in clouds: a performance evaluation. In: Jaatun, M.G., Zhao, G., Rong, C. (eds.) *Cloud Computing. LNCS*, vol. 5931, pp. 254–265. Springer, Heidelberg (2009)
32. Voorsluys, W., Broberg, J., Buyya, R.: Introduction to cloud computing. In: *Cloud Computing*, pp. 1–41 (2011)
33. Wohlin, C., Runeson, P., Höst, M., Ohlsson, M., Regnell, B., Wesslén, A.: *Experimentation in Software Engineering: An Introduction*. Kluwer Academic Publishers, Boston (2000)
34. Yazdanov, L., Fetzer, C.: Vertical scaling for prioritized VMs provisioning. In: *2012 Second International Conference on Cloud and Green Computing (CGC)*, pp. 118–125. IEEE (2012)