



Supporting CPU-based guarantees in cloud SLAs via resource-level QoS metrics

Íñigo Goiri*, Ferran Julià, J. Oriol Fitó, Mario Macías, Jordi Guitart

Barcelona Supercomputing Center and Universitat Politècnica de Catalunya, Jordi Girona 31, 08034 Barcelona, Spain

ARTICLE INFO

Article history:

Received 30 November 2010

Received in revised form

2 September 2011

Accepted 9 November 2011

Available online 25 November 2011

Keywords:

QoS metrics

SLA

Cloud provider

CPU guarantees

ABSTRACT

Success of Cloud computing requires that both customers and providers can be confident that signed Service Level Agreements (SLA) are supporting their respective business activities to their best extent. Currently used SLAs fail in providing such confidence, especially when providers outsource resources to other providers. These resource providers typically support very simple metrics like availability, or metrics that hinder an efficient exploitation of their resources.

In this paper, we propose a resource-level metric for specifying fine-grain guarantees on CPU performance. This metric allows resource providers to allocate dynamically their resources among running services depending on their demand. This is accomplished by incorporating the customer's CPU usage in the metric definition, but avoiding fake SLA violations when the customer's task does not use all its allocated resources.

We have conducted the evaluation in a virtualized provider where we have implemented the needed infrastructure for using our metric. As demonstrated in our evaluation, our solution presents fewer SLA violations than other CPU-related metrics while maintaining the Quality of Service.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

The emergence of Cloud computing [1] solutions has attracted many potential consumers, such as enterprises looking for a way to reduce costs associated with supporting their business processes. Using the Cloud, these customers can outsource services (offered by third-party service providers) which can be easily composed to build distributed applications. Additionally, the Cloud allows resource providers to offer raw resources as a service, widely known as Infrastructure as a Service (IaaS), where consumers can outsource the execution of their own services. In the same way, service providers can also use the Cloud to overcome unexpected surge demands on their services by outsourcing additional resources to other providers, or even they can fully rely on such external resources to provide their services. In all these scenarios, it is highly desirable that consumers receive fine-grain Quality of Service (QoS) guarantees from providers. This fact becomes even more essential when service providers outsource resources to resource providers, since the QoS guarantees they can provide to their respective customers depend highly on the QoS guarantees they receive from these resource providers.

Typically, a provider agrees the QoS with its customers through a Service Level Agreement (SLA), which is a bilateral contract

between both stakeholders, i.e. the customer and the provider (or between providers), which states not only the conditions and constraints of a given service, but also characterizes the agreed QoS between them using a set of metrics. Service providers naturally offer service-level metrics (e.g. service execution deadline [2], response time, throughput) to their customers for specifying the QoS. Using service-level metrics has advantages for both the customer and the provider. The former does not need to provide detailed information about the resource requirements of the service to be executed (probably the customer does not know this exactly), but only a high-level performance goal (e.g. a deadline for executing the service or a given response time threshold). The latter can freely decide the allocated resources to the service whereas it guarantees that the service meets the agreed performance goals. Being able to dynamically allocate resources to the different services is especially important for Cloud providers considering that, aiming for profitability, they tend to share their resources among multiple concurrent services owned by different customers.

On the other hand, resource providers in the Cloud must offer resource-level metrics that can be used to provide fine-grain QoS guarantees. First, the QoS agreement can be naturally expressed using resource-level metrics (e.g. number of processors, frequency of processors, FLOPS, etc.), since raw resources are the traded good. Second, having fine-grain metrics, which guarantee a given resource allocation during a time period, is especially important for service providers that outsource resources to resource providers, as we have stated before. For instance, try to figure out how a service provider could guarantee that a service will finish within a

* Corresponding author.

E-mail addresses: igoiri@ac.upc.edu (Í. Goiri), fjulia@ac.upc.edu (F. Julià), fito@ac.upc.edu (J.O. Fitó), mario@ac.upc.edu (M. Macías), jguitart@ac.upc.edu (J. Guitart).

given deadline if he does not receive fine-grain guarantees on the FLOPS supplied at every instant by the resource provider where the service is running.

Furthermore, some service providers can also benefit in some situations from supporting resource-level metrics. First, those that do not support the inference of resource requirements of a service to fulfill a given service-level metric, and for this reason, cannot offer service-level metrics in the SLA. Second, those dealing with customers (typically coming from the HPC domain) who prefer to stay with resource-level metrics that they have been using for a long time, instead of moving to service-level metrics.

Nevertheless, current Cloud providers do not support fine-grain resource-level QoS guarantees on their SLAs. In fact, most of them only support SLAs with very simple metrics based on resource availability [3–6]. For instance, whereas Amazon EC2 [7] offers different instances according to their computing capacity (which is measured in ECUs, EC2 compute units), there is not any guarantee in the SLA that this computing capacity will be supplied during the whole execution of the instance, as Amazon's SLAs only provide availability guarantees [4].

According to this, one could think on using Amazon's ECUs to support fine-grain resource-level guarantees on Cloud SLAs, considering also the porting of traditional resource-level metrics from the Grid environment to Cloud providers. However, this must be carried out carefully, since it can prevent providers from obtaining the maximum profit of their resources if done naively. In particular, metrics related to the provider's computing capacity (i.e. CPU) are especially susceptible to naive usage. For instance, if the agreed metric in the SLA establishes that the number of processors allocated to a given service must be greater or equal to some value, the provider must maintain this assignment during the whole execution of the service, even if that service is not using all the allocated capacity during some phases of its execution (i.e. the provider is forced to statically overprovision processors to the service to avoid SLA violations). Notice that the unused resources could be temporarily allocated to another service, improving in this way the utilization and profit of the provider.

In this paper, we derive a resource-level metric for specifying fine-grain QoS guarantees regarding the computing capacity of a Cloud provider by extending the Amazon's approach. Our metric overcomes the limitations of traditional CPU-related metrics. By taking into account the customer's resource usage, it allows the provider to implement dynamic resource provisioning and allocate to the different services only the amount of CPU they need. In this way, better resource utilization in the provider can be achieved. In addition, it avoids fake SLA violations, which we define as those situations where the SLA evaluator mechanism detects that an SLA is being violated, the provider will be penalized for this, but the violation is not provoked by the provider's resource allocation. In general, this occurs when the provider uses a poorly designed resource-level metric and the customer's service does not use all the resources it has allocated. Of course, the customer is free to use the amount of resources he wants, and this must not provoke any SLA violation. Therefore, the solution is to design solid resource-level metrics that support this. Finally, the proposed metric can be used in heterogeneous environments, since it is based on Amazon's ECUs.

Once the metric has been derived, we present a use case that demonstrates how it can be used in a real environment. In particular, we describe the needed infrastructure to support this metric in an SLA-driven virtualized provider, which is built using the EMO-TIVE IaaS solution [8].

The remainder of this paper is organized as follows: Section 2 explains our approach to deal with platform heterogeneity. Section 3 describes the derivation of our resource-level CPU metric. Section 4 presents the experimental environment where this

resource-level metric has been implemented. Section 5 describes the evaluation results. Section 6 presents the related work. Finally, Section 7 presents the conclusions of the paper and the future work.

2. Metric unification among heterogeneous machines

Cloud providers typically present very diverse architectures: different processor models, each of them with different clock frequencies, etc. For this reason, a good resource-level metric has to be platform-independent so it can be used in all these architectures. In this section, we describe our approach for unifying computing capacity metrics among machines with heterogeneous architectures.

Commonly, *megahertz* have been used to measure the computing capacity of a machine. Nevertheless, this does not directly measure the computing power of a machine, since noticeable differences can be observed depending on the processor architecture. For instance, using this measure an Intel Pentium III with 500 MHz would be 20 times slower than an Intel Xeon 4-core with 2.6 GHz ($\frac{42600}{500} = 20.8$). However, a simple performance comparison (e.g. *BogoMips* [9]) demonstrates that it can be up to 40 times slower ($\frac{45322.2}{499.71} = 42.6$).

In order to consider the heterogeneity of the different processor architectures, Amazon uses EC2 compute units (ECU) in its services [7]. An ECU is equivalent in CPU power to a 1.0–1.2 GHz 2007-era AMD Opteron or Intel Xeon processor. This serves as a unified measure for computing power, though it is not easily portable among different architectures. In this work, we use ECUs in order to unify CPU-related SLA metrics among heterogeneous machines, and additionally we extend Amazon's approach in order to set up SLAs that provide fine-grain CPU guarantees based on ECUs during a time period.

In our proposal, the maximum computing capacity of a given machine is measured using its maximum amount of CPU¹ (*maxCPU*) and the ECUs associated to the processor installed in that machine in the following way: $\frac{\text{maxCPU}}{100} \cdot \text{ECUs}$.

3. Derivation of CPU-based SLA metric

This section describes how our resource-level metric for establishing computing power guarantees is derived, and at the same time, discusses the limitations of alternative metrics. All the metrics discussed in this section intent to establish a guarantee on the computing performance (in terms of CPU) of a service over a time period in a given provider using the idea presented in the previous section. This guarantee is a fixed value for each SLA, represented by the SLA_i term in the formulas, which results from the negotiation between the customer and the provider. The customer is only required to specify the computing performance he requires (in terms of ECUs), which will be accepted by the provider if he is able to provide that performance.

The main difference among the metrics is how the amount of CPU for a service is defined. The more natural approach is specifying CPU performance as a function of the allocated CPU to a service, as shown in Eq. (1). This metric specifies that the computing power for a service i at every time period t has to be at least the agreed value in the SLA (SLA_i) and depends on the amount of CPU that the provider assigns to the service in that time period ($assign_i(t)$). This assignment can vary over time depending on the

¹ All CPU-related measures are quantified using the percentage of CPU used at a given time, which is the typical Linux CPU usage metric (i.e. for a computer with 4 CPUs, the maximum amount of CPU will be 400%).

provider's status, and it is periodically obtained by means of the monitoring subsystem.

$$\frac{assign_i(t)}{100} \cdot ECUs \geq SLA_i. \quad (1)$$

Note that using this metric forces the provider to statically allocate to each customer at least the amount of CPU agreed in the SLA (he can assign more if he wants), because otherwise, the SLA will be violated. Note that there is no control on whether the customer uses its allocated CPU or not. This is a quite restrictive approach, especially when customers' services have a variable CPU usage over time. In this case, the provider will suffer from low resource utilization when the services do not use all the CPU they have been allocated, since unused resources cannot be allocated to other services running in the provider.

As we have commented before, the provider aims to dynamically allocate its resources to the services depending on their demand, in order to improve resource utilization (and then increase profit). This requires an SLA metric that considers the CPU usage of services. However, as the CPU usage depends on the client's task behavior, it must be carefully used as a CPU guarantee because it can provoke undesired effects.

For instance, Eq. (2) shows an SLA metric where the computing power for a service i at every time period t depends on the amount of CPU that the service uses in that time period ($used_i(t)$). This metric assumes a provider that is able to predict the CPU usage of a given service during the next time period using the CPU usage measures of the service from the previous time periods. This could be achieved with reasonable accuracy using techniques such as the Exponential Weighted Moving Average (EWMA) [10], which uses previous observations to forecast the next value. In particular, it gives the most recent observation the greatest weight and all previous observations weights decreasing in geometric progression from the most recent back to the first. The provider will assign the CPU to the service according to its CPU usage prediction. This allows the provider to dynamically allocate the CPU among the services whereas it assigns each service at least with the CPU required to fulfill the SLA.

$$\frac{used_i(t)}{100} \cdot ECUs \geq SLA_i. \quad (2)$$

When using this metric, an SLA could be violated in two situations. First, when the provider assigns to the service an amount of CPU that is not enough to fulfill the SLA. This is a *real* violation, the provider is responsible for it, and must pay the corresponding penalty. Second, when the provider assigns to the service an amount of CPU that should be enough to fulfill the SLA, but the service does not use all the assigned CPU. This is what we have defined as a *fake* violation, since the provider is not causing it, and for this reason, he should not pay any penalty. However, Eq. (2), as currently defined, provokes this to be considered as a *real* violation, thus penalizing the provider for it.

Of course, the service should be able to use only a part of the CPU it has assigned without incurring on SLA violations. In order to allow this, we introduce our metric, in which we introduce a factor that represents the percentage of CPU used by the service with respect to its total amount of allocated CPU. As shown in Eq. (3), when the service is using all the assigned resources, Eq. (1) applies, so an SLA violation will only arise when the assigned resources are not enough to fulfill the SLA. When the service is not using all the allocated resources then the SLA is considered to be fulfilled, since the provider is not responsible that the service does not exploit its allocated resources.

$$\frac{assign_i(t)}{100} \cdot ECUs \geq SLA_i \cdot \left[\frac{used_i(t)}{assign_i(t)} \right]. \quad (3)$$

However, some services, even being CPU-intensive, do not use 100% of their assigned CPU during their whole execution. For these services, Eq. (3) does not work. Note that as these services will be considered not to use all their allocated CPU, their SLA will be always fulfilled, independently of their allocated CPU (Eq. (3) is always true). In order to overcome this limitation, we introduce an α factor as shown in Eq. (4). This factor acts as a threshold to choose when the service is considered to be using all its assigned resources. For instance, if we consider that a service is using all its allocated resources when it reaches a 90% utilization, α should be set to 0.1.

$$\frac{assign_i(t)}{100} \cdot ECUs \geq SLA_i \cdot \left[\frac{used_i(t)}{assign_i(t)} + \alpha \right]. \quad (4)$$

Operating on Eq. (4), we obtain the version of our metric ready to be used in an SLA:

$$\frac{\frac{assign_i(t)}{100} \cdot ECUs}{\left[\frac{used_i(t)}{assign_i(t)} + \alpha \right]} \geq SLA_i. \quad (5)$$

Notice that an equation in this form can have an undefined value when the denominator is zero, which happens when the service is not considered to use all its allocated resources. We avoid this by defining SLA'_i as $1/SLA_i$ and operating the equation. Notice that, when using this metric, the value specified in the SLA is SLA'_i instead of SLA_i . The final version of our metric is as follows:

$$\frac{\left[\frac{used_i(t)}{assign_i(t)} + \alpha \right]}{\frac{assign_i(t)}{100} \cdot ECUs} \leq \frac{1}{SLA_i} = SLA'_i. \quad (6)$$

4. Implementation in a virtualized resource provider

In order to demonstrate the feasibility of implementing our resource-level SLA metric, we use the SLA enforcement framework built on top of the EMOTIVE Cloud middleware [8]. EMOTIVE and the SLA framework were initially presented in [11]. EMOTIVE provides the users with virtual machines to execute the tasks, which allows taking advantage of virtualization features such as migration and easy resource management. In addition, virtualization allows the consolidation of services in the same physical resources, which enhances resource utilization. However, if the amount of resources assigned to these services is static and does not consider their real resource usage, the underutilization problem remains, as we have discussed in this paper.

The following sections describe how our metric can be implemented in top of this middleware. However, our proposal could be easily extrapolated to other implementations and used by other providers.

4.1. CPU monitoring and assignment

The resource provider has a monitoring subsystem that easily allows consulting the amount of CPU allocated to a VM ($assign_i$) and the amount that it is really using ($used_i$). The monitoring subsystem gets this information directly from the virtualization hypervisor. There are generic libraries like libvirt [12] which allow interacting with multiple hypervisors. In particular, we use the Xen Hypervisor [13] which provides all the monitoring information through the Xen API [14].

The ECUs of the underlying machine can be also calculated from the processor model, its frequency, and its associated *BogoMips* [9]. It is a direct calculation using as a basis the *BogoMips* of the Amazon EC2 reference system (i.e. 1.0–1.2 GHz 2007 Opteron or 2007 Xeon




Fig. 1. Gompertz function for SLA penalties.

processor). Then, we use these *BogoMips* to convert the computing power among different architectures into ECUs.

The amount of CPU allocated to a VM can be modified by changing the parameters of the Xen CPU scheduler policy (using the referred interfaces). By default, Xen implements the *sched-credit* policy, which works using two different parameters: capacity and weight. The capacity allows specifying the maximum amount of CPU that a VM can use. For example, if the capacity of a VM is set to 150%, it will not be able to use more than 1 CPU and a half. The weight allows defining priorities to be considered when VMs share resources. In other words, if two VMs do not have any maximum capacity restriction and the weight of the first one is two times bigger than the second one, it will be allocated with the double amount of CPU.

4.2. SLA description

The SLA framework allows assigning its own SLA to each service by using an XML description that combines both WS-Agreement [15] and WSLA [16] specifications. Using these specifications, we can accurately define the metrics derived in the previous sections. In addition, we use some features of the above-mentioned SLA specifications to define the window size of the average (i.e. 10 measures) and the interval between two consecutive measures (i.e. 2 s).

Each SLA S_i specifies the revenue that the customer will pay if the SLA is fulfilled ($Rev(S_i)$), and the penalty that the provider will pay otherwise ($Pen(S_i)$). According to this, the provider's profit for running a given application ($Prof(S_i)$) results from the revenue paid by the customer minus the penalties that the provider has to pay due to SLA violations, i.e. $Prof(S_i) = Rev(S_i) - Pen(S_i)$.

In order to establish the penalties, we use a methodology similar to the one presented in [17], which is built over the same Cloud infrastructure. Each penalty $Pen(S_i)$ is calculated as a percentage of the revenue obtained when fulfilling the corresponding SLA in the following way: $Pen(S_i) = Rev(S_i) \cdot \frac{Gom(\sigma(S_i))}{100}$. This percentage is calculated by using a Gompertz function [18], which is a kind of sigmoid function. Its basic form is $y(t) = ae^{be^{ct}}$, where a is the upper asymptote, b sets the x displacement, c sets the growth rate, and e is Euler's Number. Note that b and c are negative numbers.

For our purposes, we have adapted this function as shown in Fig. 1, which displays the penalty percentage depending on a $\sigma(S_i)$ function that represents the SLA violation ratio. In particular, as shown in Eq. (7), when this ratio is zero or less, the penalty percentage is 0. When this ratio is one, the percentage tends to the maximum penalty percentage (MAX_p). MAX_p for SLA S_i is defined as a percentage of the price that the client pays for that SLA. Notice that this upper limit ($MAX_p\%$ of the price) can be agreed with the client during the SLA negotiation process.

$$Gom(\sigma(S_i)) = \begin{cases} 0 & \text{if } \sigma(S_i) \leq 0 \\ MAX_p \cdot e^{-e^{-7 \cdot \sigma(S_i) + e}} & \text{otherwise.} \end{cases} \quad (7)$$

As shown in Eq. (8), $\sigma(S_i)$ function depends on the number of violations occurring for that SLA (V_i) and is parameterized with two




Fig. 2. SLA evaluation cycle.

thresholds, $\lambda_{\min}$ and $\lambda_{\max}$. These thresholds indicate the minimum number of SLA violations in order to start paying penalties and the maximum number of SLA violations that could occur during the execution of the application, respectively.

$$\sigma(S_i) = \frac{V_i}{\lambda_{\max} - \lambda_{\min}} - \frac{\lambda_{\min}}{\lambda_{\max} - \lambda_{\min}}. \quad (8)$$

4.3. SLA evaluation

Fig. 2 shows the SLA evaluation cycle implemented in the virtualized provider. The *Direct Measurer* component gets the values from a monitoring subsystem, which is implemented using daemons running in the Xen Domain-0. There is one daemon per VM, which obtains all the resource usage information of this VM via XML-RPC calls to the Xen API. The *Direct Measurer* controls the measurement intervals for each metric in the SLA and ensures the refresh of the measures on correct time. When these values arrive at the *Measurer Sched* component, it checks if any metric has updated its value. In this case, the *Measurer Sched* recalculates the top-level metric defined in the SLA using the corresponding formula (*Recalculate SLA*) and compares the result with the agreed value specified in the guarantee terms of the SLA (SLA_i) (*Evaluate SLA*). If the SLA is fulfilled, the *Measurer Sched* waits until the next iteration, otherwise an SLA violation notification is raised.

Note that two different entities use this SLA evaluation procedure. On one side, the virtualized provider has an internal SLA enforcement mechanism, whose main goal is to allow the provider to detect ongoing SLA violations and initiate resource management actions in order to solve such violation (for instance, allocating more resources to the VM). This SLA enforcement is executed every second and it is transparent to the users. On the other side, there is a SLA fulfillment entity that evaluates the SLAs in order to check whether the terms agreed in each SLA between the customer and the provider are being fulfilled or not. In our experiments, this SLA fulfillment is executed every 100 s and it is what determines the penalties the provider will pay to the users.

5. Evaluation

The use case presented in the previous section is used in our evaluation as the testbed to compare how our metric performs with respect to other CPU-based metrics introduced in Section 3. In the first part, we set up the SLA framework in a resource provider using the EMOTIVE middleware [8] and we evaluate the different metrics using real executions with a small proof-of-concept workload composed of three tasks in order to show in




Fig. 3. Comparison of CPU-based SLA metrics. (a) CPU assignment and usage for Task 1, 2, and 3. (b) Fulfillment of SLA metric A for Task 1. (c) Fulfillment of SLA metric B for Task 1. (d) Fulfillment of SLA metric C for Task 1.

detail the consequences of using each metric. In the second part, we use a simulator that mimics the operation of the resource provider (including the same SLA framework) in order to be able to evaluate the different SLA metrics with a real workload in a large scale environment.

5.1. Proof-of-concept workload

We have conducted three experiments. Each of them consists of the concurrent execution of three tasks, which are basically CPU-consuming tasks with variable CPU requirements over time, in a single node of the virtualized provider during 400 s. Each task runs within a dedicated virtual machine. This makes the system able to dynamically change the resources allocated to each virtual machine and allows demonstrating the benefits of our metric in such a scenario. This node is a 64-bit architecture with 4 Intel Xeon CPUs at 2.6 GHz and 16 GB of RAM, which runs Xen 3.3.1 [13] in the Domain-0. This quad-core processor has been measured to have 21288.8 (4*5322.2) BogoMips, which correspond to 10.4 ECUs approximately. Each task is started at a different time to simulate what could happen in a real scenario, where different customers can reach the provider at different times.

Each task has its own SLA, which describes the agreed QoS with the provider. For each experiment, a different SLA metric is used to specify this QoS. In particular, in the first experiment all tasks use the metric referred in Eq. (1) (from now on denoted as *SLA Metric A*), in the second one they use the metric referred in Eq. (2) (denoted as *SLA Metric B*), and in the last one they use the metric referred in Eq. (6) (denoted as *SLA Metric C*).

For each experiment, we have monitored the allocated CPU to each task, its real CPU usage, and the value of the SLA metric at each moment in time. Comparing this value with the guarantee agreed in the SLA, we can know if the SLA is being violated or not. Fig. 3 displays all this information. The top graphic in this figure shows the CPU assignment and usage for each task over time. These values will be the same independently of the metric used in the SLA. We have forced the provider to execute the three tasks, but there are not enough resources for fulfilling all the SLAs. This allows us to evaluate how the different metrics deal with real SLA violations. The provider distributes the resources in such a way that the SLAs of *Task2* and *Task3* are never violated in these experiments, while *Task1* behavior shows all the possible situations regarding SLA management. For this reason, we focus the explanation only on

Task1, which has negotiated an SLA with a CPU requirement of 6 ECUs (1/6 for Metric C).

As shown in the top graphic of Fig. 3, at the beginning of the experiment *Task1* is assigned with the whole machine, since it is the only task running in the provider. When *Task2* arrives at the provider at second 150, the CPU is redistributed among *Task1* and *Task2* according to their requirements. Finally, at second 250, *Task3* arrives at the provider, which redistributes again the CPU among the three tasks. Regarding the CPU usage of *Task1*, it starts its execution consuming almost 400% of CPU. After 100 s of execution, it reduces its CPU consumption and starts using just 20% of CPU for the next 40 s. Then, it begins increasing its CPU usage until it consumes around 150% of CPU.

Next three graphics in Fig. 3 show the value of the SLA metric over time and its agreed value in the SLA, for Metric A, B and C, respectively. As described in Section 3, *SLA Metric A* only takes into account the amount of resources assigned to the task. While this value is above of the agreed value in the SLA, the SLA is fulfilled. Otherwise, an SLA violation arises. As shown in the second graphic, due to the redistribution of CPU occurred when *Task3* arrives at the provider at second 250, the value of the *SLA Metric A* for *Task1* falls below the agreed value in the SLA, thus violating it. The shaded zone identifies in the figure the interval where *Task1* is violating the SLA (from second 250 to 400). In order to avoid this, the provider would be forced to refuse executing *Task3*, even if this would be initially possible since the other tasks are not using all their allocated CPU. Notice that using this metric makes the provider unable to exploit adequately its resources, in the sense that it should always assign enough resources to fulfill the SLA, whether the tasks use them or not.

On the other side, *SLA Metric B* considers the resource usage of the tasks. Using this metric allows the provider to freely move the CPU assignment if tasks do not use all their allocated resources. Nevertheless, as shown in the third graphic, this metric can give fake SLA violations. In particular, when the CPU usage of *Task1* goes below the agreed value in the SLA (from second 115 to 400), the SLA is being violated. However, not all these violations are a responsibility of the provider. In particular, between second 115 and 250 (light gray zone in the graphic), the SLA is being violated because although the provider is giving enough resources to guarantee the SLA, the task does not use all its assigned CPU. As we have described before, this is a fake violation. On the other side, from second 250 to 400, the SLA is being violated (dark gray

Table 1
Duration (in s) of SLA violations (proof-of-concept workload).

SLA metric	Real	Fake	Total
A	150	0	150
B	150	135	285
C	110	0	110

zone in the graphic) because the provider has not allocated enough resources to the task to fulfill its SLA, and consequently, even if it used all of them, its CPU usage would be under the SLA threshold. This is then a real violation. Notice that the value of *SLA Metric B* has some start delay with respect to the CPU usage. This is because of the different sampling rates between the SLA monitoring system and the CPU usage monitoring system.

Finally, *SLA Metric C* (the one proposed in this paper) gets the best from the two other ones, since it considers both the CPU assigned to a task and its real CPU usage. As discussed in Section 3, this metric specifies the inverse value in the SLA with respect to the other metrics. For this reason, when using this metric, an SLA violation will happen when the measured value of the metric is greater than the agreed one (1/6 in this particular experiment). As shown in the fourth graphic, this only occurs in the shaded zone, that is, from second 290 to 400. As commented before, this is the only zone where the task is using all its allocated resources and the provider has not allocated enough resources to the task to fulfill its SLA. On the other side, notice that when using this metric fake SLA violations are avoided, even when the task is not using all its allocated resources (as occurs from second 115 to 290). In these experiments, we have considered that a task is using all its allocated resources when it uses more than 70% of them (this corresponds to a 0.3 value for the α factor in Eq. (6)).

Summarizing, three well-differentiated phases can be appreciated in the figure. In the first one, the assigned CPU is enough to fulfill the SLA and *Task1* is using all its allocated resources. Consequently, none of the metrics is violating the SLA. In the second one, CPU allocated to *Task1* has been reduced, but the task is not using all its allocated resources. In this case, using the *SLA Metric A* causes low resource utilization because unused resources cannot be assigned to any other task without violating the SLA, while using the *SLA Metric B* causes fake SLA violations. Finally, in the last phase, the resource assignment is not enough to fulfill the SLA, and *Task1* uses all its assigned resources. In this case, all the metrics violate the SLA.

The time while the SLA is being violated for each metric is summarized in Table 1. It shows the time (in s) with real SLA violations, with fake violations, and the total time violating the SLA. Notice that, when using *SLA Metric C*, the time with real SLA violations is lower than when using the other metrics. This occurs because *SLA Metric C* takes into account that *Task1* is not using all its allocated resources between second 250 and second 290. In addition, *SLA Metric B* is the only metric causing fake SLA violations.

5.2. Real workload

In this section, we evaluate the different metrics by executing a real workload. The workload is a whole day load obtained from Grid5000 [19] corresponding to Monday first of October, 2007. 215 jobs of 10 different applications are submitted during this period. The size of our real testbed is not enough to execute this workload. For this reason, we have built a simulated provider composed of 24 nodes, which mimics the behavior of the virtualized provider used in previous experiments. The number of nodes has been chosen in order to be able to execute the maximum number of concurrent jobs of the workload, while still violating some SLAs to demonstrate the behavior of the different metrics.

The simulator reads the workload and starts tasks when the time to do it comes. This simulator mimics the actual behavior of EMOTIVE and considers the times to create VMs and start running tasks. When the task is started, the SLA framework is notified and receives the XML describing the SLA for this task. At this point, the simulator interacts with the actual SLA framework as it was the actual system providing the information of the system status periodically.

We set up three types of nodes in the simulated provider in order to demonstrate how our proposal can deal with heterogeneous machines. The first type is the one already used in previous experiments. The other two nodes types comprise a 64-bit architecture with 4 Intel Xeon CPUs at 3 GHz and 16 GB of RAM, and a 64-bit architecture with 8 Intel Xeon CPUs at 2.8 GHz and 16 GB of RAM, which respectively correspond to 12 and 22.4 ECUs. For each task, we specify a static requirement of resources, and the dynamic resource usage it performs during its execution. Each task is executed within a VM with 2 VCPUs, since this is enough to fulfill the CPU requirements of all the applications. In order to demonstrate the impact on each metric of having dynamic CPU usage pattern over time, we have forced two of the 10 applications to use more CPU than the requested during some parts of their execution. Resource requirements and usages constitute the input of the simulator, which implements a backfilling policy taking into account these requirements in order to decide the allocated resources to each task and simulates the specified resource consumption for each task in a given time slot.

Each task has a different SLA. All of them use the same metric for specifying the CPU guarantee, but with a different agreed value, which depends on the task CPU requirement. In addition, the SLA specifies a revenue of \$0.38 per hour (we use the pricing of an Amazon EC2 medium instance [7], which has similar features to the VMs used in the experiment). The SLA is evaluated every 100 s, and the penalty is calculated following Eq. (7) using a maximum penalty (MAX_p) of 125%. Note that this corresponds to the SLA fulfillment procedure described in Section 4.3.

As the SLA is evaluated every 100 s, the maximum number of SLA violations ($\lambda_{\max}$) occurs when the SLA is violated in every sample. Hence, $\lambda_{\max}$ is equal to the number of samples, which is calculated by dividing the application execution time (T_{exec}) by the sampling period (T_{sample}). Using this information, and considering that $\lambda_{\min}$ is zero in our case, we can operate on Eq. (8) and calculate $\sigma(S_i)$ as follows:

$$\sigma(S_i) = \frac{V_i \cdot T_{\text{sample}}}{T_{\text{exec}}}.$$

Table 2 shows the obtained results using each one of the metrics, including the number of violations (*real*, *fake*, and *total*) and the final profit. For *SLA Metric C*, we have also evaluated the impact of the α factor. These results demonstrate that using our metric the provider can get the highest profit, since it avoids all the fake SLA violations and reduces the real ones when the application is not using all the allocated resources. In particular, we have reduced the incurred penalties in more than a 58% regarding *SLA Metric A* and more than a 72% compared to *SLA Metric B* when $\alpha = 0.1$. In addition, note that the lower the alpha factor is, the lower the number of SLA violations is, though the difference is not very noticeable. Nevertheless, this difference highly depends on the CPU usage of the applications. This occurs because for higher alpha values the task is considered to be using all its allocated resources during more time, which increases the probability of SLA violations imputable to the provider.

These results allow us to conclude that our metric has globally less SLA violations than the other ones, it is able to avoid fake SLA violations, and it allows the provider to freely redistribute the resources when the tasks are not using them.

Table 2
Number of SLA violations and profit (real workload).

SLA metric	(α)	Real	Fake	Total	Profit (\$)
A		3891	0	3891	231.0
B		3891	2857	6748	176.6
C	0.05	1674	0	1674	277.5
C	0.1	1782	0	1782	271.3
C	0.2	1877	0	1877	266.9
C	0.4	1884	0	1884	266.6

6. Related work

As Cloud technology is still in its early stage, current Cloud providers only support the specification of QoS guarantees by means of SLAs in a very simple way. For instance, the SLAs supported by Amazon EC2 [7] only consider an availability metric, the so-called *annual uptime percentage* [4]. It is noticeable that violations must be monitored and claimed by the user. Furthermore, whereas Amazon EC2 offers different instances according to their ECUs, there is no guarantee in the SLA that this computing capacity will be supplied during the whole execution of the instance. Similarly, GoGrid also considers availability metric [5], including the availability of the storage and the primary DNS. It also offers some metrics related with network performance such as jitter, latency, and packet loss rate. Analogously, 3Tera also provides availability guarantees on its Virtual Private Datacenter [3] in the same way RackSpace does [6]. Nevertheless, none of the Cloud approaches for SLA considers CPU-related metrics since they implement static provisioning of resources. As this can lead to low resource utilization, moving to dynamic provisioning is recommended to increase profit.

Besides, it is important to emphasize the advantages of Cloud Computing when managing the heterogeneity. Conventional legacy Utility computing systems toolkits, such as *Autopilot* [20], implemented a set of resource management policies that acted in conjunction with distributed sensors that provide real-time performance data. Based on these, a set of actuators implemented policy decisions for increasing throughput in heterogeneous workloads. However, Cloud computing brings some features that can greatly simplify the process of server consolidation. In fact, many other Cloud Computing features, such as runtime migration of VMs, could be used in addition to strengthen the consolidation.

Until now, SLAs have been primarily used to provide QoS guarantees on e-business computing utilities. Due to this, most of the works focusing in this area [21–25] have used metrics of interest in these environments such as throughput, response time, and availability. None of them has included resource-level metrics as part of its solution, as we suggest in this paper.

In addition, recent contributions around the Cloud computing paradigm are focused on business-level SLAs. [26] states that security measures in SLAs are needed to establish trust between clients and providers; [27] exploits a high-level metric, i.e. Quality of Experience (QoE), in order to allow customers to evaluate providers' services; and [28] determines that SLA offers in Cloud environments must be oriented on the business perspective, e.g. trust, privacy, and security requirements. However, none of these proposals considers what should be the basis of these high-level SLAs, which is the QoS guarantee based on accurate fine-grain metrics, such as CPU usage, in order to meet users' expectations.

On the other side, some effort has been carried out in the Grid environment in order to provide QoS guarantees using SLAs. Some proposals focus on service-level metrics. For instance, [29,30] propose SLA management systems for Grid environments which use the expected completion time of the jobs for defining SLA metrics. Similarly, [31] defines latency, throughput, and availability metrics and tries to minimize execution times and costs in a

Grid provider. Besides, Kyriazis et al. [32] present an innovative QoS-aware workflow mapping mechanism in Grid environments. Basically, it takes into account the offered QoS level, in terms of availability and execution time, when mapping workflow processes to server instances.

Whereas these last previous works neglect resource-level metrics in their solution, other works define metrics for providing CPU guarantees in the SLA, but usually in a very limiting way for the resource provider. For instance, [33] allows defining the CPU count, the CPU architecture, and the CPU performance when executing a physics application on the Grid. Similarly, [34] uses metrics such as the number of processors, the CPU share, and the amount of memory. As we have discussed, when used incorrectly, some of these metrics force the provider to statically overprovision CPU to the applications in order to avoid SLA violations, while others induce fake SLA violations when the applications do not use all their allocated resources. As demonstrated in the evaluation, our solution overcomes these limitations.

7. Conclusion

In this paper, we have proposed a resource-level SLA metric for specifying fine-grain QoS guarantees regarding the computing capacity of a Cloud provider. This metric overcomes the limitations of traditional CPU-related metrics. In particular, it takes into account if services are really using their allocated resources. This allows the provider to implement dynamic resource provisioning, allocating to the different services only the amount of CPU they need. In this way, better resource utilization in the provider can be achieved. Furthermore, dynamic resource provisioning can be accomplished while avoiding fake SLA violations that could happen with other metrics when customers' services do not use all the resources they have allocated. Finally, it can be used in heterogeneous machines, since it is based on Amazon's ECUs. As demonstrated in the evaluation, using this metric reduces the number of SLA violations, since it only detects those that are a real responsibility of the provider.

Although we have derived a metric for managing CPU, the idea of considering the real resource usage of services can be used with other resources such as disk I/O, memory, or network. This is part of our future work. In addition, we are working on mechanisms for supporting the inference of resource requirements of a service to fulfill a given service-level metric.


Acknowledgments

This work is supported by the Ministry of Science and Technology of Spain and the European Union (FEDER funds) under contract TIN2007-60625 and grant AP2008-0264, by the Generalitat de Catalunya (contract 2009-SGR-980), and by the European Commission under FP7-ICT-2009-5 contract 257115 (OPTIMIS).

References

- [1] R. Buyya, C.S. Yeo, S. Venugopal, J. Broberg, I. Brandic, Cloud computing and emerging it platforms: vision, hype, and reality for delivering computing as the 5th utility, *Future Generation Computer Systems* 25 (2009) 599–616. doi:10.1016/j.future.2008.12.001. URL: <http://portal.acm.org/citation.cfm?id=1528937.1529211>.
- [2] Aneka Software, 2010, Website: <http://www.manjrasoft.com/products.html>.
- [3] 3Tera SLA, 2010, Website: <http://www.3tera.com/News/Press-Releases/Recent/3Tera-Introduces-the-First-Five-Nines-Cloud-Computing.php>.
- [4] Amazon EC2 Service Level Agreement, 2010, Website: <http://aws.amazon.com/ec2-sla/>.
- [5] GoGrid Service Level Agreement, 2010, Website: <http://www.gogrid.com/legal/sla.php>.
- [6] Rackspace Cloud Servers Service Level Agreement, 2010, Website: http://cloudservers.rackspacecloud.com/index.php/Cloud_Servers_SLA.
- [7] Amazon EC2, 2010, Website: <http://aws.amazon.com/ec2/>.
- [8] EMOTIVE Cloud, 2010, Website: <http://www.emotivecloud.net>.


- [9] W. van Dorst, BogoMips mini-Howto, Tech. Rep., Clifton Scientific Text Services V38, March 2006.
- [10] S. Roberts, Control chart tests based on geometric moving averages, *Technometrics* 1 (3) (1959) 239–250.
- [11] I. Goiri, F. Julia, J. Ejarque, M. Palol, R. Badia, J. Guitart, J. Torres, Introducing virtual execution environments for application lifecycle management and SLA-driven resource distribution within service providers, in: 8th IEEE International Symposium on Network Computing and Applications, NCA'09, 2009, pp. 211–218.
- [12] libvirt: the virtualization API, Website, 2010. <http://libvirt.org/>.
- [13] Xen Hypervisor, 2010. Website: <http://www.xen.org>.
- [14] E. Mellor, R. Sharp, D. Scott, Xen management API, revision 1.0.6. July 2008, Website, 2010.
- [15] A. Andrieux, K. Czajkowski, A. Dan, K. Keahey, H. Ludwig, J. Pruyne, J. Rofrano, S. Tuecke, M. Xu, Web services agreement specification (WS-agreement), in: Global Grid Forum GRAAP-WG, Draft, 2004.
- [16] A. Keller, H. Ludwig, The WSLA framework: specifying and monitoring service level agreements for web services, *Journal of Network and Systems Management* 11 (1) (2003) 57–81.
- [17] J.O. Fitó, I. Goiri, J. Guitart, SLA-driven elastic cloud hosting provider, in: 18th Euromicro Conference on Parallel, Distributed and Network-Based Processing, PDP'10, Pisa, Italy, 2010, pp. 111–118.
- [18] D. Ratkowsky, *Handbook of Nonlinear Regression Models*, M. Dekker, 1990.
- [19] The grid workloads archive, 2010. Website: <http://gwa.ewi.tudelft.nl>.
- [20] R.L. Ribler, H. Simitci, D.A. Reed, The autopilot performance-directed adaptive control system, *Future Generation Computer Systems* 18 (2001) 175–187. doi: 10.1016/S0167-739X(01)00051-6. URL: <http://portal.acm.org/citation.cfm?id=514474.514489>.
- [21] B. Abrahao, V. Almeida, J. Almeida, A. Zhang, D. Beyer, F. Safai, Self-adaptive SLA-driven capacity management for Internet services, in: 10th IEEE/IFIP Network Operations and Management Symposium, NOMS 2006, Vancouver, Canada, 2006, pp. 557–568.
- [22] K. Appleby, S. Fakhouri, L. Fong, G. Goldszmidt, S. Krishnakumar, D. Pazel, J. Pershing, B. Rochwerger, Oceano—SLA-based management of a computing utility, in: Symposium on Integrated Network Management, IM'01, Seattle, WA, USA, 2001, pp. 855–868.
- [23] M. Bennani, D. Menasce, Resource allocation for autonomic data centers using analytic performance models, in: 2nd International Conference on Autonomic Computing, ICAC'05, Seattle, USA, 2005, pp. 229–240.
- [24] G. Lodi, F. Panzieri, D. Rossi, E. Turrini, SLA-driven clustering of QoS-aware application servers, *IEEE Transactions on Software Engineering* 33 (3) (2007) 186–197.
- [25] J. Xu, M. Zhao, J. Fortes, R. Carpenter, M. Yousif, On the use of fuzzy modeling in virtualized data center management, in: 4th International Conference on Autonomic Computing, ICAC'07, Jacksonville, USA, 2007, p. 25.
- [26] F. Dickmann, M. Brodhun, J. Falkner, T. Knoch, U. Sax, Technology transfer of dynamic IT outsourcing requires security measures in SLAs, in: Proceedings of the 7th International Workshop on Grid Economics and Business Models, GECON 2010, Ischia, Italy, 2010, pp. 1–15.
- [27] K. Tserpes, F. Aisopos, D. Kyriazis, T. Varvarigou, Service selection decision support in the Internet of services, in: Proceedings of the 7th International Workshop on Grid Economics and Business Models, GECON 2010, Ischia, Italy, 2010, pp. 16–33.
- [28] I. Ul-Haq, I. Brandic, E. Schikuta, SLA validation in layered cloud infrastructures, in: Proceedings of the 7th International Workshop on Grid Economics and Business Models, GECON 2010, Ischia, Italy, 2010, pp. 153–164.
- [29] A. Ching, L. Sacks, P. McKee, SLA management and resource modelling for grid computing, in: London Communications Symposium, LCS'03, London, UK, 2003.
- [30] A. Sahai, S. Graupner, V. Machiraju, A. van Moorsel, Specifying and monitoring guarantees in commercial grids through SLA, in: 3rd IEEE/ACM International Symposium on Cluster Computing and the Grid, CCGrid 2003, Tokyo, Japan, 2003, pp. 292–299.
- [31] D. Menasce, E. Casalicchio, Quality of service aspects and metrics in grid computing, in: Computer Measurement Group Conference, CMG'04, Las Vegas, USA, 2004, pp. 521–532.
- [32] D. Kyriazis, K. Tserpes, I. Papagiannis, K. Konstanteli, T. Varvarigou, A QoS-based selection mechanism exploiting business relationships in workflows, in: Proceedings of the 6th International Workshop on Grid Economics and Business Models, GECON 2009, Springer, Delft, Netherlands, 2009, pp. 102–115.
- [33] L. Skital, M. Janusz, R. Slota, J. Kitowski, Service level agreement metrics for real-time application on the grid, *Lecture Notes in Computer Science* 4967 (2008) 798–806.
- [34] H. Raj, R. Nathuji, A. Singh, P. England, Resource management for isolation enhanced cloud services, in: 2009 ACM Cloud Computing Security Workshop, CCSW'09, ACM, Chicago, IL, USA, 2009, pp. 77–84. <http://doi.acm.org/10.1145/1655008.1655019>.



Íñigo Goiri obtained his bachelors degree in computer science from the University of Lleida in 2006 and his Masters degree in Computer Architecture and Network Systems (CANS) from the Technical University of Catalonia (UPC) where he has also started his Ph.D. studies. He is currently working at Barcelona Supercomputing Center (BSC) as researcher staff in the eDragon Research Group.



Ferran Julià has a masters degree in Physics at the University of Barcelona (UB) and a bachelors degree in Computer Science at the Technical University of Catalonia (UPC), where he is currently finishing a Masters degree in Computer Architecture and Network Systems (CANS). In addition, he works as a researcher at Barcelona Supercomputing Center (BSC). His research interests are simulation and power efficient computing in cloud systems.



J. Oriol Fitó is a Ph.D. student in Computer Science at the Computer Architecture Department of the Technical University of Catalonia (UPC). He obtained his Bachelors degree in computer science and Masters degree in Computer Architectures, Networks and Systems (CANS) from the UPC in 2007 and 2009, respectively. Moreover, he is currently a research engineer at the Barcelona Supercomputing Center (BSC).



Mario Macías obtained his Masters Degree in Computer Sciences at the Universitat Autònoma de Barcelona. After some years working in IT companies, he joined Barcelona Supercomputing Center (BSC) in 2006 to work in European Research Projects and carry out his Ph.D. studies at Technical University of Catalonia.



Jordi Guitart received his MS and Ph.D. degrees in Computer Science from the Technical University of Catalonia (UPC), in 1999 and 2005, respectively. Currently, he is a collaborator professor at the Computer Architecture Department of the UPC and an associate researcher at Barcelona Supercomputing Center (BSC). His research interests are oriented towards innovative resource management approaches for modern distributed computing systems. He is involved in a number of European projects.