

COMPUTING SCIENCE

Consumer Side Resource Accounting in the Cloud

Ahmed Mihoob, Carlos Molina-Jimenez and Santosh Shrivastava

TECHNICAL REPORT SERIES

No. CS-TR-1259

June 2011

Consumer Side Resource Accounting in the Cloud

A. Mihoob, C. Molina-Jimenez and S. Shrivastava

Abstract

The paper investigates the possibility of consumer side resource accounting of cloud services, meaning, whether it is possible for a consumer to independently collect all the resource usage data required for calculating billing charges for pay-per-use cloud services. The accounting models of two widely used cloud services are examined and possible sources of difficulties are identified, including causes that could lead to discrepancies between the metering data collected by the consumer and the provider. The investigation is motivated by the fact that cloud service providers perform their own measurements to collect usage data, but as yet there are no equivalent facilities of consumer-trusted metering that are commonly provided by utility service providers like gas and electricity. The paper goes on to suggest how cloud service providers can improve their accounting models to facilitate the task of consumer side resource accounting.

Bibliographical details

MIHOOB, A., MOLINA-JIMENEZ, C. AND SHRIVASTAVA, S.

Uncertainty and Uncertainty Tolerance in Service Provisioning

[By] A. Mihoob, C. Molina-Jimenez and S. Shrivastava

Newcastle upon Tyne: Newcastle University: Computing Science, 2011.

(Newcastle University, Computing Science, Technical Report Series, No. CS-TR-1259)

Added entries

NEWCASTLE UNIVERSITY

Computing Science. Technical Report Series. CS-TR-1259

Abstract

The paper investigates the possibility of consumer side resource accounting of cloud services, meaning, whether it is possible for a consumer to independently collect all the resource usage data required for calculating billing charges for pay-per-use cloud services. The accounting models of two widely used cloud services are examined and possible sources of difficulties are identified, including causes that could lead to discrepancies between the metering data collected by the consumer and the provider. The investigation is motivated by the fact that cloud service providers perform their own measurements to collect usage data, but as yet there are no equivalent facilities of consumer-trusted metering that are commonly provided by utility service providers like gas and electricity. The paper goes on to suggest how cloud service providers can improve their accounting models to facilitate the task of consumer side resource accounting.

About the authors

Ahmed Mihoob received his BSc in Computer Science from Sebha University- Libya in 1991, and his MSc in Computing Science from School of Computing Science - University of Newcastle in 2003. He is currently working towards his PhD under supervision of Prof. Santosh Shrivastava. His broad area of research interest is distributed computing, currently, he is particularly interested in exploring pricing and billing of cloud services. Cloud computing allows consumers to perform computation in a public cloud with a pricing scheme typically based on incurred resource consumption. Consumers should be able to configure their applications to minimise charges. This in turn requires them to monitor and measure resource consumption to enable run time configuration management of applications. This is a relatively new research area that examines cloud computing from a new angle.

Dr. Carlos Molina-Jimenez received his PhD in computing science from University of Newcastle in 2000 for work on anonymous interactions in the Internet. Currently he is a Research Associate at the School of Computing Science, University of Newcastle where he is a member of the Distributed Systems Research Group. He is working on the EPSRC funded research project on Networked Computing in Inter-organisational Settings where he has been responsible for developing the architectural concepts of virtual organisations, trust management and electronic contracting.

Santosh Shrivastava was appointed a Professor of Computing Science, University of Newcastle upon Tyne in 1986; he leads the Distributed Systems Research Group. He received his Ph.D. in computing science from Cambridge in 1975. His research interests are in the areas of distributed systems, fault tolerance and application of transaction and workflow technologies to e-commerce, virtual organisations and Grid-based systems. His group is well known as the developers of an innovative distributed transaction system, called Arjuna and a CORBA based dependable workflow system for the Internet. He has led and managed several European Union funded, multipartner research projects, beginning with BROADCAST (1992) to a project on complex service provisioning on the Internet, TAPAS, that started in 2002. Together with his colleagues he set up a company in 1998 in Newcastle to productise Arjuna transaction and workflow technologies. Now based within the University campus, Arjuna Technologies is a centre of excellence in transaction technologies and is focusing on building products to support reliable Web Services-based applications. Close industry-university collaboration is guaranteed, and research projects on E-commerce platform and services have been initiated.

Suggested keywords

CLOUD RESOURCE CONSUMPTION
STORAGE AND COMPUTATIONAL RESOURCES
RESOURCE METERING AND ACCOUNTING MODELS
AMAZON WEB SERVICES

Consumer Side Resource Accounting in the Cloud

Ahmed Mihoob, Carlos Molina–Jimenez and Santosh Shrivastava

School of Computing Science
Newcastle University
Newcastle upon Tyne, NE1 7RU, UK
{a.m.mihoob, carlos.molina, santosh.shrivastava}@ncl.ac.uk

Abstract. Keywords: cloud resource consumption, storage and computational resources, resource metering and accounting models, Amazon Web Services

1 Introduction

Cloud computing services made available to consumers range from providing basic computational resources such as storage and compute power (infrastructure as a service, IaaS) to sophisticated enterprise application services (software as a service SaaS). A common business model is to charge consumers on a pay-per-use basis where they periodically pay for the resources they have consumed. Needless to say that for each pay-per-use service, consumers should be provided with an unambiguous resource accounting model that precisely describes all the constituent chargeable resources of the service and how billing charges are calculated from the resource usage (resource consumption) data collected on behalf of the consumer over a given period. If the consumers have access to such resource usage data then they can use it in many interesting ways, such as, making their applications billing aware, IT budget planning, create brokering services that automate the selection of services in line with user's needs and so forth. Indeed, it is in the interest of the service providers to make resource consumption data available to consumers; incidentally all the providers that we know of do make such data accessible to their consumers in a timely fashion.

An important issue that is raised is the *accountability* of the resource usage data: who performs the measurement to collect the resource usage data - the provider, the consumer, a trusted third party (TTP), or some combination of them¹? The traditional utility providers such as water, gas and electricity perform their own measurements to collect usage data using metering devices (trusted by consumers) that are deployed in the consumers premises. Cloud service providers also perform their own measurements to collect usage data,

¹ A note on terminology: 'accountability' refers to concepts such as responsibility, answerability, trustworthiness; not to be confused with 'resource accounting' that refers to the process concerned with calculating financial charges.

although, as yet there are no equivalent facilities of consumer-trusted metering; rather, consumers have no choice but to take whatever usage data that is made available by the provider as trustworthy. A good introduction into this issue can be found in [16].

In light of the above discussion, it is worth investigating whether it is possible for a consumer (or a TTP acting on behalf of the consumer) to independently collect all the resource usage data required for calculating billing charges. In effect, this means that a consumer (or a TTP) should be in a position to run their own metering service for measuring resource consumption. If this were possible, then consumers will be able to perform reasonableness checks on the resource usage data available from service providers as well as raise alarms when apparent discrepancies are suspected in consumption figures; furthermore, innovative charging schemes can be constructed with confidence by consumers who are themselves offering third party brokering services. In our earlier paper [15], we developed the notion of a *Consumer-centric Resource Accounting Model* for a cloud service. We say that a resource accounting model is *weakly consumer-centric* if all the data that the model requires for calculating billing charges can be queried programmatically from the provider. Further, we say that an accounting model is *strongly consumer-centric* if all the data that the model requires for calculating billing charges can be collected independently by the consumer (or a TTP). Strongly consumer-centric accounting models have the desirable property of openness and transparency, since service users are in a position to verify the charges billed to them. In that paper we also evaluated the accounting model of a given cloud infrastructure service (simple storage service, S3 from Amazon) to see how well it matches the proposed notion. We independently collected (by examination of requests and responses) our own resource usage data for S3 and compared it with the provider's data. Our investigations indicated that even though it is a very simple service, the accounting model description nevertheless has a few ambiguities and not all the data that the model requires for calculating billing charges can be queried programmatically from the provider. However, we pointed out that most of these shortcomings can be fixed fairly easily by Amazon to make their model weakly as well as strongly consumer-centric. This paper contributes to the prior work in three ways:

- the evaluation of accounting models is extended to include a compute service (Amazon Elastic Compute Cloud, EC2) and we point out a few ambiguities in the EC2 model description (Section 3);
- we precisely identify the causes that could lead to discrepancies between the metering data collected by the provider and the consumer, and explain how the discrepancies can be resolved (Section 4); since Amazon accounting models are representative of the IaaS cloud providers, our ideas are applicable generally; and,
- we present ideas on how an accounting model should be constructed so as to make them strongly consumer-centric (Section 5).

2 Background

For resource accounting it is necessary to determine the amount of resources consumed by a given consumer (also called client and subscriber) during given time interval, for example, a billing period. Accounting is routinely performed by providers with the help of accounting systems, to produce records to bill the consumer at the end of the billing period. **Accounting systems** are composed of three basic services: **metering**, **accounting** and **billing**.

We show a consumer side accounting system in Fig.1.

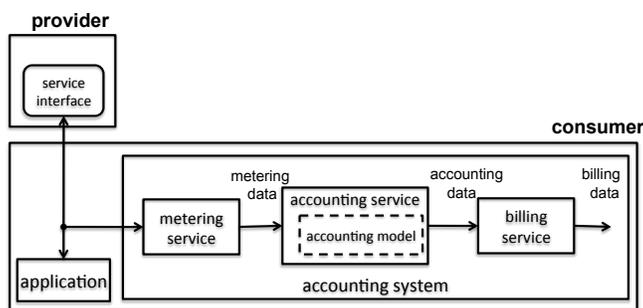


Fig. 1. Consumer side resource accounting system

We assume that resources are exposed as services through one or more service interfaces. As shown in the figure, the metering service intercepts the message traffic between the consumer application and the cloud services and extracts relevant data required for calculating resource usage (for example, the message size which would be required for calculating bandwidth usage). The metering service stores the collected data for use by the accounting service. The accounting service retrieves the metering data, computes resource consumption from the data using its *accounting model* and generates accounting data that is needed by the billing service to calculate the billing data.

Accounting models are provider-specific in the sense that the functionality of an accounting model is determined by the provider's policies. These policies determine how the metrics produced by his metering service are to be interpreted; for example, 1.7 GB of storage consumption can be interpreted by the provider's accounting model either as 1 or 2 GB.

The accounting models of cloud providers are normally available from their web pages and in principle can be used by subscriber to perform their own resource accounting. The difficulty here for the subscriber is to extract the accounting model from their online documentation as most providers that we know of, unnecessarily blur their accounting models with metering and billing parameters.

The parameters involved in accounting models depend on the type of service (SaaS, PaaS, IaaS, etc.) offered. In this paper we will examine, from the point of view of consumer side resource accounting, the accounting models of Amazon's Simple Storage Service (S3) and Elastic Compute Cloud (EC2). In the following discussion, we gloss over the fine details of pricing, but concentrate on metering and accounting services.

3 Accounting of Resource Consumption

3.1 S3 Accounting Model

An S3 space is organised as a collection of buckets which are similar to folders, except that they do not support nesting. A bucket can contain zero or more objects of up to 5 terabytes of data each; an object is simply a file uploaded by the customer from their local disk into their S3 space. Both buckets and objects are identified by names (keys in Amazon terminology) chosen by the customer. S3 provides SOAP and RESTful interfaces. An S3 customer is charged for: a) **storage**: storage space consumed by the objects that they store in S3; b) **bandwidth**: network traffic generated by the operations that the customer executes against the S3 interface; and c) **operations**: number of operations that the customer executes against the S3 interface.

Storage: The key parameter in calculation of the storage bill is number of byte hours accounted to the customer. *Byte Hours* (ByteHrs) is the the number of bytes that a customer stores in their account for a given number of hours.

Amazon explains that *the GB of storage billed in a month is the average storage used throughout the month. This includes all object data and metadata stored in buckets that you created under your account. We measure your usage in TimedStorage-ByteHrs, which are added up at the end of the month to generate your monthly charges.* They further state that *at least twice a day, we check to see how much storage is used by all your Amazon S3 buckets. The result is multiplied by the amount of time passed since the last checkpoint.*

From the definition of ByteHrs it follows that to calculate their bill, a customer needs to understand 1) how their byte consumption is measured, that is, how the data and metadata that is uploaded is mapped into consumed bytes in S3; and 2) how Amazon determines the number of hours that a given piece of data was stored in S3 —this issue is directly related to the notion of a checkpoint.

Amazon explains that each object in S3 has, in addition to its data, system metadata and user metadata; furthermore it explains that the **system metadata** is generated and used by S3, whereas **user metadata** is defined and used only by the user and limited to 2 KB of size [1]. Unfortunately, Amazon does not explain how to calculate the actual storage space taken by data and metadata. To clarify the issue, we conducted a number of experiments (see [15]) involving uploading of a number of objects of different names, data and user metadata into an equal number of empty buckets.

Three conclusions can be drawn from these experiments: first, the mapping between bytes uploaded (as measured by intercepting upload requests) and bytes stored in S3 correspond one to one; second, the storage space occupied by system metadata is the sum of the lengths (in Bytes) of object and bucket names and incur storage consumption; third, user metadata does not impact storage consumption. In summary, for a given uploaded object, the consumer can accurately measure the total number of bytes that will be used for calculating ByteHrs.

Next, we need to measure the 'Hrs' of 'ByteHrs'. As stated earlier, Amazon states that at least twice a day they check the amount of storage consumed by a customer. However, Amazon does not stipulate exactly when the checkpoints take place.

To clarify the situation, we conducted a number of experiments that consisted in uploading to and deleting files from S3 and studying the Usage Reports of our account to detect when the impact of the PUT and DELETE operations were accounted by Amazon. Our findings are summarised in Fig.2. It seems that, currently, Amazon does not actually check customers' storage consumption twice a day as they specify in their Calculating Your Bill document, but only once. From our observations, it emerged that the time of the checkpoint is decided randomly by Amazon within the 00:00:00Z and 23:59:59Z time interval².

In the figure, CP stands for checkpoint, thus $CP_{30} : 2GB$ indicate that CP_{30} was conducted on the 30th day of the month at the time specified by the arrow and reported that at that time the customer had 2 GB stored in S3. SC stands for Storage Consumption and is explained below.

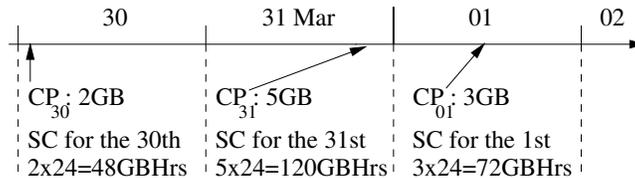


Fig. 2. Amazon's checkpoints.

As shown in the figure, Amazon uses the results produced by a checkpoint of a given day, to account the customer for the 24 hrs of that day, regardless of the operations that the customer might perform during the time left between the checkpoint and the 23:59:59Z hours of the day. For example, the storage consumption for the 30th will be taken as $2 \times 24 = 48$ GBHrs; where 2 represents the 2GB that the customer uploaded on the 30th and 24 represents the 24 hrs of the day.

² S3 servers are synchronised to the Universal Time Coordinated (UTC) which is also known as the Zulu Time (Z time) and in practice equivalent to the Greenwich Mean Time (GMT)

Bandwidth: Amazon charges for the network data transferred from the customer to S3 ('DataTransfer-In', the request message of an operation) and the network data transferred from S3 to the customer (the corresponding response message, 'DataTransfer-Out'). It is however not clear from the available information how the size of the message is calculated. We therefore conducted several experiments involving uploading, downloading, deleting etc. of objects using both RESTful and SOAP interfaces and compared the information extracted from the intercepted messages with the information available from Amazon usage reports. It turns out that for RESTful operations, only the size of the object (in DataTransfer-In for PUT, and DataTransfer-Out for GET) is taken into account and system and user metadata is not part of the overhead, whereas for SOAP operations, the total size of the message is taken into account.

Operations: It is straightforward for a consumer to count the type and number of operations performed on S3. A simple charging scheme is used: a block of 1000 operations composed of PUT, COPY, POST or LIST costs one cent, whereas a block of 10000 GET and all other operations, excluding DELETE, costs one cent. Delete operations are free. We note that an operation might fail to complete successfully. The error response in general contains information that helps identify the party responsible for the failure: the customer or the S3 infrastructure. For example, *NoSuckBucket* errors are caused by the customer when they try to upload a file into a non-existent bucket; whereas an *InternalError* code indicates that S3 is experiencing internal problems. Our understanding is that the consumer is charged for an operation, whether the operation succeeded or not.

3.2 EC2 Accounting Model

EC2 is a computation service offered by Amazon as an IaaS [4]. The service offers raw virtual CPUs (also called a Virtual Machine or VM) to subscribers. A subscriber is granted administrative privileges over his VM, that he can exercise by means of sending remote commands to the Amazon Cloud from his desktop computer. For example, he is expected to configure, launch, stop, re-launch, terminate, backup, etc. his VM. In return, the subscriber is free to choose the operating system (eg Windows or Linux) and applications to run.

In EC2 terminology, a running virtual CPU is called an *instance* whereas the frozen bundle of software on disk that contains the libraries, applications and initial configuration settings that are used to launch an instance is called the *Amazon Machine Image (AMI)*.

Currently, Amazon offers six types of instances that differ from each other in four initial configuration parameters that cannot be changed at running time: amount of EC2 compute units that it delivers, size of their memory and local storage (also called ephemeral and instance storage) and the type of platform (32 or 64 bits). An EC2 compute unit is an Amazon unit and is defined as the equivalent CPU capacity of a 1.0–1.2 GHz 2007 Opteron or 2007 Xeon processor.

Thus Amazon offer small, large, extra large and other types of instances. For example, the default instance type is the *Small Instance* and is a 32 bit platform that delivers 1 EC2 compute unit and provided with 1.7 GB of memory and 160 GB of local storage.

These types of instances are offered to subscribers under several billing models: **on-demand instances**, **reserved instances** and **spot instances**. In our discussion we will focus on on-demand instances.

Under the on-demand billing model, Amazon defines the unit of consumption of an instance as the *instance hour* (*instanceHrs*). Currently, the cost of an instance hour of a small instance running Linux or Windows, is, respectively, 9.5 and 12 cents. On top of charges for instance hours, instance subscribers normally incur additional charges for data transfer that the instances generates (Data Transfer In and Data Transfer Out) and for additional infrastructure that the instance might need such as disk storage, IP addresses, monitoring facilities and others. As these additional charges are accounted and billed separately, we will leave them out of our discussion and focus only on instance hours charges.

The figures above imply that if a subscriber accrues 10 instanceHrs of a small instance consumption, running Linux, during a month, he will incur a charge of 95 cents at the end of the month.

In principle, the pricing tables publicly available from Amazon web pages should allow a subscriber to independently conduct his own accounting of EC2 consumption. In the absence of a well defined accounting model this is not a trivial exercise.

Insights into the EC2 accounting model are spread over several on-line documents from Amazon. Some insight into the definition of instance hour is provided in the *Amazon EC2 Pricing* document [3] (see just below the table of *On-demand Instances*) where it is stated that *Pricing is per instance-hour consumed for each instance, from the time an instance is launched until it is terminated. Each partial instance-hour consumed will be billed as a full hour.* This statement suggests that once an instance is launched it will incur at least an instance hours of consumption. For example, if the instance runs continuously for 5 minutes, it will incur 1 instanceHrs; likewise, if the instance runs continuously for 90 minutes, it will incur 2 instanceHrs.

The problem with this definition is that it does not clarify when an instance is considered to be launched and terminated. Additional information about this issue is provided in the *Billing* section of FAQs [2], *Paying for What You Use* of the *Amazon Elastic Compute (Amazon EC2)* document [4] and in the *How You're Charged* section of the User Guide [5]. For example, in [4] it is stated that *Each instance will store its actual launch time. Thereafter, each instance will charge for its hours of execution at the beginning of each hour relative to the time it launched.*

From information extracted from the documents cited above it is clear that Amazon starts and stops counting instance hours as the instance is driven by the subscriber, through different states. Also, it is clear that Amazon instance hours are accrued from the execution of one or more individual sessions executed

(possibly empty) of operations that Amazon executes when the event occurs, to count the numbers of instance hours consumed by the instance.

There are two types of events: subscriber's and internal to the FSM events. The subscriber's events are the commands (*launch*, *application commands*, *reboot*, *stop* and *terminate*) that the subscribers issues to operate his instance; likewise, internal events are events that occur independently from the subscriber's commands, namely, *timer = 60min* and *failure*.

AMI configured: is the initial state. It is reached when the subscriber successfully configures his AMI so that it is ready to be launched.

running: is the state where the instance can perform useful computation for the subscriber, for example, it can respond to application commands issued by the subscriber.

terminated: is the final state and represents the end of the life cycle of the instance. Once this state is reached the instance is destroyed. To perform additional computation after entering this state the subscriber needs to configure another AMI. The terminated state is reached when the subscriber issues the *terminate* command, the instance fails when it was in running state or the instance fails to reach running state.

shuttingdown: is reached when the subscriber issues the *reboot* command.

stopped: this state is supported only EBS-backed instances (S3-backed instances cannot be stopped) and is reached when the user issues *stop* command, say for example, to perform backup duties.

States and instance hours: In the figure, *NinstHrs* is used to count the number of instance hours consumed by an instance during a single session. The number of instance hours consumed by an instance is determined by the integer value stored in *NinstHrs* when the instance reaches the *terminated* state. *timer* is Amazon's timer to count 60 minutes intervals; it can be set to zero (*timer = 0*) and started (*starttimer*).

In the FSM, the charging operations are executed as suggested by the Amazon's on line documentation. For example, in *Paying for What You Use* Section of [4], Amazon states that the beginning of an instance hour is relative to the launch time. Consequently, the FSM sets *NinstHrs = 1* when the subscriber executes a launch command from the *AMI configured* state. At the same time, *timer* is set to zero and started. *NinstHrs = 1* indicates that once a subscriber executes a launch command, he will incur at least one instance hour. If the subscriber leaves his instance in the *running* state for 60 minutes (*timer = 60min*) the FSM increments *NinstHrs* by one, sets the timer to zero and starts it again. From *running* state the timer is set to zero when the subscriber decides to terminate his instance (*terminate* command) or when the instance fails (*failure* event). Although Amazon's documentation does not discuss it, we believe that the possibility of an instance not reaching the *running* state cannot be ignore, therefore we have included a transition from *pending* to *terminated* state; the FSM sets the timer to zero when this abnormal event occurs.

As explained in *Basics of Amazon ESB-Backed AMIs and Instances and How You're Charged* of [5], a running ESB-backed instance can be stopped by the subscriber by means of the *stop* command and drive it to the *stopped* state. As indicated by *timer = 0* operation executed when the subscriber issues a *stop* command, an instance in *stopped* state incurs no instance hours. However, though it is not shown in the figure as this is a different issue, Amazon charges for EBS storage and other additional services related to the stopped instance. The subscriber can drive an instance from the *stopped* to the *terminated* state. Alternatively he can re-launch his instance. In fact, the subscriber can launch, stop and launch his instance as many times as he needs to. However, as indicated by the *NinstHrs ++*, *timer = 0* and *starttimer* operations over the arrow, every transition from *stopped* to *pending* state accrues an instance hours of consumption, irrespectively of the time elapsed between each pair of consecutive *launch* commands.

Experiments with Amazon instances: To verify that the accounting model described by the FSM of Fig. 3-a) matches Amazon's description, we (as subscribers) conducted a series of practical experiments. In particular, our aim was to verify how the number of instance hours is counted by Amazon.

The experiments involved 1) configuration of different AMIs; 2) launch of instances; 3) execution of remote commands to drive the instances through the different states shown in the FSM. For example, we configured AMIs, launched and run them for periods of different lengths and terminated them. Likewise, we launched instances and terminated them as soon as they reached the *running* state.

To calculate the number of instance hours consumed by the instances, we recorded the time of execution of the remote commands *launch*, *stop*, *terminate* and *reboot*, and the time of reaching both transient and permanent states. For comparison, we collected data (start and end time of an instance hour, and number of instance hours consumed) from Amazon EC2 usage report.

A comparison of data collected from our experiments against Amazon's data from their usage report reveals that currently, the beginning of an instance hour is not the execution time of the subscriber's *launch* command, as documented by Amazon, but the time when the instance reaches the *running* state. These findings imply that the current accounting model currently in use is the one described by the FSM of Fig. 3-b). As shown in the figure, the *NinstHrs* is incremented when the instance reaches the *running* state.

4 Potential Causes of Discrepancies

4.1 Storage

Since, for the calculation of *ByteHrs*, the time of the checkpoint is decided randomly by Amazon within the 00:00:00Z and 23:59:59Z time interval, the time

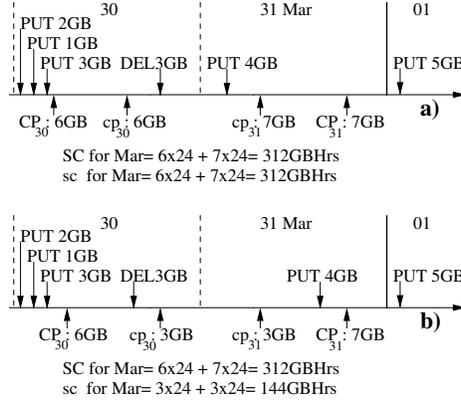


Fig. 4. Impact of checkpoints.

used at the consumer’s side need not match that at the provider’s side: a potential cause for discrepancy. This is illustrated with the help of Fig.4.

The figure shows the execution time of four PUT and one DEL operations executed by an S3 consumer during the last two days of March. The first day of April is also shown for completeness. For simplicity, the figure assumes that the earliest PUT operation is the very first executed by the consumer after opening his S3 account. The figure also shows the specific points in time when checkpoints are conducted independently by two parties, namely, Amazon and a consumer. Thus, CP and cp represent, respectively, Amazon’s and the consumer’s checkpoints; the Giga Bytes shown next to CP and cp indicate the storage consumption detected by the checkpoint. For example, on the 30th, Amazon conducted its checkpoint about five in the morning and detected that, at that time, the customer had 6 GB stored ($CP_{30} : 6GB$). On the same day, the consumer conducted his checkpoint just after midday and detected that, at that time, he had 6 GB stored ($cp_{30} : 6GB$). SC and sc represent, respectively, the storage consumption for the month of March, calculated by Amazon and consumer, based on their checkpoints.

The figure demonstrates that the storage consumption calculated by Amazon and consumer might differ significantly depending on the number and nature of the operations conducted within the time interval determined by the two parties’ checkpoints, for example, within CP_{31} and cp_{31} .

Scenario a) shows an ideal situation where no consumer’s operations are executed within the pair of checkpoints conducted on the 30th or 31st. The result is that both parties calculate equal storage consumptions. In contrast, b) shows a worse–case scenario where the DEL operation is missed by CP_{30} and counted by cp_{30} and the PUT operation is missed by cp_{31} and counted by CP_{31} ; the result of this is that Amazon and the consumer, calculate SC and sc, respectively, as 312 GB and 144 GB.

Ideally, Amazon’s checkpoint times should be made known to consumers to prevent any such errors. Providing this information for upcoming checkpoints is perhaps not a sensible option for a storage provider, as the information could be ‘misused’ by a consumer by placing deletes and puts around the checkpoints in a manner that artificially reduces the consumption figures. An alternative would be to make the times of past checkpoints available (e.g., by releasing them the next day).

Impact of network and operation latencies: In the previous discussion concerning calculation of ByteHrs (illustrated using Fig. 4), we have implicitly assumed that the execution of a PUT (respectively a DELETE) operation is an atomic event whose time of occurrence is either less or greater than the checkpoint time (i.e., the operation happens either before or after the checkpoint). This allowed us to say that if the checkpoint time used at the provider is known to the consumer, then the consumer can match the ByteHrs figures of the provider. However, this assumption is over simplifying the distributed nature of the PUT (respectively a DELETE) operation. In Fig.5 we explicitly show network and operation execution latencies for a given operation, say PUT; also, i , j , k and l are provider side checkpoint times used for illustration. Assume that at the provider side, only the completed operations are taken into account for the calculation of ByteHrs; so a checkpoint taken at time i or j will not include the PUT operation (PUT has not yet completed), whereas a checkpoint taken a time k or l will. What happens at the consumer side will depend on which event (sending of the request or reception of the response) is taken to represent the occurrence of PUT. If the timestamp of the request message (PUT) is regarded as the time of occurrence of PUT, then the consumer side ByteHrs calculation for a checkpoint at time i or j will include the PUT operation, a discrepancy since the provider did not! On the other hand, if the timestamp of the response message is regarded as the time of occurrence of PUT, then a checkpoint at time k will not include the PUT operation (whereas the provider has), again a discrepancy. In short, for the operations that occur ‘sufficiently close’ to the checkpoint time, there is no guarantee that they get ordered identically at both the sides with respect to the checkpoint time.

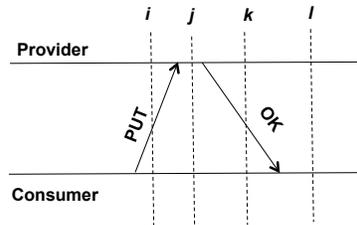


Fig. 5. Network and operation latencies

Operations: Earlier we stated that it is straightforward for a consumer to count the type and number of operations performed on S3. There is a potential for discrepancy caused by network latency: operations that are invoked 'sufficiently close' to the end of an accounting period (say i) and counted by the consumer for that period, might get counted as performed in the next period (say j) by the provider if due to the latency, these invocation messages arrive in period j . This will lead to the accumulated charges for the two period not being the same. This is actually not an issue, as the Amazon uses the timestamp of the invocation message for resolution, so the consumer can match the provider's figure.

One likely source of difficulty about the charges for operations is determining the liable party for failed operations. Currently, this decision is taken unilaterally by Amazon. In this regard, we anticipate two potential sources of conflicts: DNS and propagation delays. As explained by Amazon, some requests might fail and produce a Temporary Redirect (HTTP code 307 error) due to temporary routing errors which are caused by the use of alternative DNS names and request redirection techniques [17]. Amazon's advice is to design applications that can handle redirect errors, for example, by resending a request after receiving a 307 code (see [1], Request Routing section). Strictly speaking these errors are not caused by the customer as the 307 code suggests. It is not clear to us who bears the cost of the re-tried operations.

4.2 EC2

The mismatch between Amazon's documented accounting model and the one currently in use (Fig. 3-a and b, respectively) might result in discrepancies between the subscriber's and Amazon's calculations of instance hours. For example, imagine that it takes five minutes to reach the *running* state. Now imagine that the subscriber launches an instance, leaves it running for 57 minutes and then terminates it. The subscriber's *NinstHours* will be equal to two: $NinstHours = 1$ at launch time and then $NinstHours$ is incremented when $timer = 60min$. In contrast, to the subscriber's satisfaction, Amazon's usage records will show only one instance hour of consumption. One can argue that this discrepancy is not of the subscriber's concern since, economically, it always favours him.

More challenging and closer to the subscriber's concern are discrepancies caused by failures. Amazon's documentation does not stipulates how instances that fail accrue instance hours. For example, examine Fig. 3-a) and imagine that an instance suddenly crashes after spending 2 hrs and 15 min in *running* state. It is not clear to us whether Amazon will charge for the last 15 min of the execution as a whole instance hour. As a second example, imagine that after being launched either from *AMI configured* or *stopped* states, an instance progresses to *pending* state and from there, due to a failure, to *terminated*. It is not clear to us if Amazon will charge for the last instance hour counted by *NinstHours*.

We believe that, apart from these omissions about failure situations, the accounting model of Fig. 3-a) can be implemented and used by the subscriber to produce accurate accounting. A salient feature of this model is that all the events

(*launch*, *stop* and *terminate*) that impact the *NinstHrs* counter are generated by subscriber. The only exception is the *timer = 60min* event, but that can be visible to the subscriber if he synchronises his clock to UTC time.

The accounting model that Amazon actually uses (Fig. 3–b) is not impacted by failures of instances to reach *running* state because in this model, *NinstHrs* is incremented when the instance reaches *running* state. However, this model is harder for the subscriber to implement since the event that causes the instance to progress from *pending* to *running* state is not under the subscriber’s control.

5 Related Work

Interest in resource accounting has been expressed in different application domains (old time-sharing systems, grid computing, Internet traffic, phone networks, cloud computing, etc.) where a limited amount of resources are offered to the public. Consequently, and as pointed out in [20], different communities use term *accounting* to refer to different aspects of the billing process. For instance, some authors use the term to refer to the all embracing process of metering, collecting, interpreting and reporting costing and charging-related information of the usage of a service or resource, while others use it to refer to only one of the subprocesses. In our paper, as explained in Section 2, we use the term *accounting service* to refer exclusively to the process that computes the metering data provided by the metering service and *accounting system* to refer to the whole process.

No consensus has been reached yet about cloud computing terminology. Terms like client, user, subscriber and consumer mean different things to different authors. A good attempt to remedy the situation is [6]. For instance, the draft provides an illuminating discussion about what SaaS, PaaS and IaaS should offer to be categorised as such and what they charge for.

An early work on resource accounting in time-sharing systems is presented in [14]; the authors discuss how different pricing and charging mechanisms can help in cost recovery, resource allocation and regulation of demand for computing resources.

In our paper we focus on accounting of resources offered under usage-based charges, which is an alternative to flat-rate based charging mechanisms. The pros and cons of the two alternatives are discussed in [18]. The author argues in favour of the simplicity and efficiency of the flat-rate based mechanisms.

Since the focus of our paper is on accounting services, we leave out of the discussion the metering service which is directly related to accounting. Metering is a topic in its own right and a challenging one when in addition to the traditional time and volume metrics, it counts for QoS measurement and trust-related concerns. A good introduction into metering is presented in [19].

An architecture for accounting and billing for resources consumed in a federated Grid infrastructure is suggested in [10]. The paper provides a valuable insight into the requirements (resource re-deployment, SLA awareness, pre-paid and post-paid billing, standardised records and others) that accounting and

billing services should meet. In [7], the author discuss similar requirements for accounting and billing services, but within the context of federated network of telecommunication providers. Both papers overlook the need to provide consumers with means of performing consumer-side accounting. A detailed discussion of an accounting system similar to the one shown in Fig. 1 but aimed at telecommunication services is provided in [12].

In [11] the authors present a taxonomy of billing models and an illuminating discussion about the metering parameters (e.g. volume consumed, start and end time of a session) that each model requires.

Accounting models are fundamental to subscribers interested in planning for minimisation of expenditures on cloud resources. The questions raised are what workload to outsource, to which provider, what resources to rent, when, and so on. Examples of research results in this direction are reported in [24, 8]. In [8] the authors discuss how an accounting service deployed within an organisation can be used to control expenditures on public cloud resources; their accounting service relies on data downloaded from the cloud provider instead of calculating it locally.

In [13], the authors present a framework to help subscribers choose between two or more providers, considering performance and monetary costs of accessing their main services. They conducted a number of experiments against six of the leading cloud providers, including Amazon. One of their experiments aimed at evaluating the monetary costs of virtual machine usage. Unfortunately the discussion of the accounting model used is not elaborated.

In [9], the authors take Amazon cloud as an example of cloud provider and estimate the performance and monetary-cost to compute a data-intensive (terabytes) workflow that requires hours of CPU time. The study is analytical (as opposite to experimental) and based on the authors' accounting model. For instance, to produce actual CPU-hours, they ignore the granularity of Amazon instance hours and assume CPU seconds of computation. This work stresses the relevance of accounting models.

Amazon, like other providers, charges for instance hours which represent anything between one second and 60 minutes of instance running time; arguably, this granularity might be too high and inconvenient when one tries to make users (say employees within a company) accountable for the actual hours of VM time (as opposite to instance hours) that they consume from a public cloud. This question is raised in [21]. In the same order, it is important to stress that a 60 min run does not necessarily mean 60 min of actual processor time; some authors [23] have observed that Amazon EC2 small instances typically receive only a 40% to 50% share of the processor.

Currently, Amazon charges for the number of operations executed against its S3 interface and independently of the internal computation cost that the request generates. This issue has been studied in [22].

6 Concluding Remarks

We investigated whether it is possible for a consumer (or a TTP acting on behalf of the consumer) to independently collect, for a given cloud service, all the resource usage data required for calculating billing charges. We examined two main IaaS services: storage and compute from Amazon; our investigation revealed the causes that could lead to discrepancies between the metering data collected by the the consumer not matching that of the provider. Essentially these causes can be classed into three categories discussed below.

1. Incompleteness and ambiguities: We pointed out several cases where an accounting model specification was ambiguous or not complete. For example, regarding bandwidth consumption, it is not clear from the available information what constitutes the size of a message. It is only through experiments we worked out that for RESTful operations, only the size of the object (in DataTransfer-In for PUT, and DataTransfer-Out for GET) is taken into account and system and user metadata is not part of the message size, whereas for SOAP operations, the total size of the message is taken into account. Failure handling is another area where there is lack of information and/or clarity: for example, concerning EC2, it is not clear how instances that fail accrue instance hours.
2. Unobservable events: If an accounting model uses one or more events that impact resource consumption, but these events are not observable to (or their occurrence cannot be deduced accurately by) the consumer, then the data collected at the consumer side could differ from the that of the provider. Calculation of storage consumption in S3 (ByteHrs) is a good example: here, the checkpoint event is not observable.
3. Differences in the measurement process: Difference can arise if the two sides use different techniques for data collection. Calculation of ByteHrs again serves as a good example. We expect that for a checkpoint, the provider will directly measure the storage space actually occupied, whereas, assuming the checkpoint time is known, the consumer will mimic the process by adding (for PUT) and subtracting (for DELETE) to calculate the space, and as we discussed with respect to fig.5, discrepancies are possible.

References

1. Amazon: Amazon simple storage service. developer guide, API version 2006-03-01 (2006), www.amazon.com
2. Amazon: Amazon ec2 faqs (2011), aws.amazon.com/ec2/faqs
3. Amazon: Amazon ec2 pricing (2011), aws.amazon.com/ec2/pricing
4. Amazon: Amazon elastic compute cloud (amazon ec2) (2011), aws.amazon.com/EC2/
5. Amazon: Amazon elastic compute cloud user guide (api version 2011-02-28) (2011), docs.amazonwebservices.com/AWSEC2/latest/UserGuide/

6. Badger, L., Grance, T., Patt-Corner, R., Voas, J.: Draft cloud computing synopsis and recommendations. Tech. Rep. NIST Special Publication 800-146, National Institute of Standards and Technology (NIST) (May 2011), available at: <ftp://ftp.isi.edu/in-notes/rfc2828.txt>
7. Bhushan, B., Tschichholz, M., Leray, E., Donnelly, W.: Federated accounting: Service charging and billing in a business-to-business environment. In: Proc. 2001 IEEE/IFIP Int'l Symposium on Integrated Network Management VII. pp. 107–121 (2001)
8. den Bossche, R.V., Vanmechelen, K., Broeckhove, J.: Cost-optimal scheduling in hybrid iaas clouds for deadline constrained workloads. In: Proc. IEEE 3rd Int'l Conf. on Cloud Computing(Cloud'10). pp. 228–235 (2010)
9. Deelman, E., Singh, G., Livny, M., Berriman, B., Good, J.: The cost of doing science on the cloud: The montage example. In: Proc. Int'l Conf. on High Performance Computing, Networking, Storage and Analysis (SC'08) (2008)
10. Elmroth, E., Marquez, F.G., Henriksson, D., Ferrera, D.P.: Accounting and billing for federated cloud infrastructures. In: The Eighth Int'l Conf. on Grid and Cooperative Computing. pp. 268–275. Aug. 27–28, Lanzhou, Gansu, China (2009)
11. Kim, S.D., Park, S.H., Keum, C., Chung, T.M.: A study new challenge for billing system in converged service platform. In: Proc. 6th Int'l Conf. on Networked Computing and Advanced Information Management (NCM'10). pp. 390–395 (2010)
12. de Leostar, E., McGibney, J.: Flexible multi-service telecommunications accounting system. In: Proc. Int'l Network Conf. (INC'00) (2000)
13. Li, A., Yangi, X., Zhang, S.K.M.: Cloudcmp: Shopping for a cloud made easy (2010)
14. McKell, L.J., Hansen, J.V., Heitger, L.E.: Charging for computing resources. *Computing Surveys* 11(2), 105–120 (Jun 1979)
15. Mihoob, A., Molina-Jimenez, C., Shrivastava, S.: A case for consumer-centric resource accounting models. In: Proc. IEEE 3rd Int'l Conf. on Cloud Computing(Cloud'10). pp. 506–512 (2010)
16. Molina-Jimenez, C., Cook, N., Shrivastava, S.: On the feasibility of bilaterally agreed accounting of resource consumption. In: 1st Int'l Workshop on Enabling Service Business Ecosystems (ESBE08). pp. 170–283. Sydney, Australia (2008)
17. Murty, J.: *Programming Amazon Web Services*. O'Reilly (2008)
18. Odlyzko, A.: Internet pricing and the history of communications (Feb 2001), <http://www.research.att.com/~amo>
19. Pias, M., Wilbur, S., Bhatti, S., Crowcroft, J.: Securing the internet metering and billing. In: Proc. IEEE Global Telecommunications Conf. (GLOBECOM'02). pp. 1603–1607 (2002)
20. Pras, A., van Beijnum, B.J., Sprenkels, R., Parhonyi, R.: Internet accounting. *IEEE Communications Magazine* 39(5), 110–113 (May 2001)
21. Vecchiola, C., Pandey, S., Buyya, R.: High-performance cloud computing: A view of scientific applications. In: Proc. 10th Int'l Symposium on Pervasive Systems, Algorithms and Networks (I-SPAN'09) (2009)
22. Wachs, M., Xu, L., Kanevsky, A., Ganger, G.R.: Exertion-based billing for cloud storage access. In: Proc. 3rd USENIX Workshop on Hot Topics in Cloud Computing (HotCloud'11) (2011)
23. Wang, G., Ng, T.S.E.: The impact of virtualization on network performance of amazon ec2 data center. In: Proc. 29th IEEE Conf. on Computer Communications (INFOCOM'10). pp. 1–9 (2010)

24. Wang, H., Jing, Q., Chen, R., He, B., Qian, Z., Zhou, L.: Distributed systems meet economics: Pricing in the cloud. In: Proc. 2nd USENIX Workshop on Hot Topics in Cloud Computing (HotCloud'10) (2010)