An Open Market of Cloud Data Services

Verena Kantere
Institute of Services Science, University of Geneva, Geneva, Switzerland

Keywords: Cloud Computing Services, Cloud Data Services.

Abstract: Cloud data services are a very attractive solution for the management of business and large-scale data. The incapability of creating fully functional data services and, moreover, providing them with clear guarantees on technical and business aspects, creates a problematic cloud service-provisioning situation. We propose the development of a framework that will enable the systematic and efficient creation and management of cloud data services. Such a framework is necessary in order to achieve the exchange of cloud data services in any open market, where cloud providers and their customers can advertise the offered and requested data services in a free manner and make contracts for service provisioning. We discuss the skeleton of such a framework, which comprises three parts: first, profiling data services, second, designing service offers and demands, and, search and match of data services.

1 INTRODUCTION

Cloud computing has given rise to a new perspective of data management that adds more optimization dimensions, beyond traditional performance, such as cost, availability, elasticity, scalability etc. In a cloud environment this is realized with the provisioning of *data services*, which includes the transparent management of infrastructure, data and workload. Cloud data services are a very attractive solution for the management of business and large-scale data. Even though the cloud industry offers today data services, these mainly concern key-value stores and lack substantial functionality of a traditional DBMS (e.g., support for transactions, optimization techniques, fully fledged declarative query language).

The incapability of creating thorough data services and, moreover, providing them with clear guarantees on technical and business aspects, creates a problematic cloud service-provisioning situation. On one hand, cloud providers cannot design efficient and appropriate data services and, on the other, consumers are not aware of their data management needs. Moreover, providers cannot describe accurately their offered data services and, thus, consumers do not have enough information on their received data services. This results in misuse of data services for long periods and expensive migration to new data services. Therefore, even though cloud data services are a critical component of cloud-hosted applications, their deficiencies prohibit their wide applicability.

We argue that we need to fill the gap between providers and consumers of data services that exists in today's cloud business, by not only solving the above two issues, but also offering an all-inclusive solution for the provision of efficient and appropriate data services. We propose the development of a framework that will enable the systematic and efficient creation and management of cloud data services. Such a framework is necessary in order to achieve the exchange of cloud data services in any open market, where cloud providers and their customers can advertise the offered and requested data services in a free manner and make contracts for service provisioning. The openness of such a market eliminates the imbalance of the two parties, by allowing everybody to contribute and request services and enabling the near-real time negotiation of data service provision. The proposed framework comprises three parts: first, a novel technique for profiling data services; second, a methodology and methods for designing service offers and demands; and, finally, methods and algorithms for search and match of data services.

Such a framework will directly benefit cloud business and industry by enabling them to: first, design fully functional data services tailored to the needs of the customer, and, second, give visibility

Kantere V..
An Open Market of Cloud Data Services.
DOI: 10.5220/0004955004390444
In *Proceedings of the 4th International Conference on Cloud Computing and Services Science* (CLOSER-2014), pages 439-444
ISBN: 978-989-758-019-2

and clarity to offer and demand, which will lead to optimal service provisioning. Potentially, the research results on data service search and match will give tools for new practices around service engineering and trading.

In the following we discuss how the proposed framework relates to and advances the domain of cloud data service provisioning.

2 CLOUD DATA SERVICES

The new trend for service infrastructures in the IT domain, cloud computing, is also a new area of research (Armbrust, M., et al. 2010). Recently, the research community has shown an enormous interest in the area of cloud data management. The reason is that cloud computing seems to be the ideal paradigm for the time- and cost- efficient management of big amounts of data, such as scientific data of large-scale. Beyond scientific data, cloud data management can be applicable to other data, such as personal data, commercial data, and any kind of archives, public or private. Users (persons, companies, organizations) may store their data in the cloud, which provides data management services. The cloud alleviates the burden of data management from the users for remuneration.

Emerged cloud data service provisions such as the Amazon SimpleDB[1], Amazon Relational Database Service[2], Microsoft Azure[3] and Google Cloud SQL[4] offer preliminary data management services, but are a long way from offering full capabilities of a traditional database management system (DBMS). Moreover, cloud systems that support massively distributed data management, such as Amazon Dynamo (DeCandia, G., et al. 2007), Google Bigtable (Chang, F. et al. 2006), MicroSoft SQL Server for cloud (Bernstein, P. A. et al. 2011), Yahoo PNUTS[5], Cassandra[6] and HBase[7] are concerned with data consistency and availability issues for analytical data in key-value formats, leaving transactional management on relational data out of their scope.

Yet, the business and enterprise world is rapidly turning towards the cloud data-serving systems to cover their data management needs in a cheap and easy, but also, efficient and reliable, fashion. Such data management applications require high levels of functionality, i.e. the functionality of a traditional DBMS, seamlessly offered through a cloud provider. Therefore, it is absolutely necessary to evolve the first generation of cloud data-serving systems to complete data management systems. The proposed framework is an output towards this direction, by enabling the formulation and the provision of fully functional data services through the cloud.

Current research projects in the general area of cloud data management deal with issues of data consistency and manipulation (Kraska, T. et al. 2009), and issues of data filtering and aggregation[8]. Before even dealing with such issues, several research groups, (e.g. (Abadi, D. J., 2009, Aboulnaga, A. et al. 2009)), discuss the benefits, drawbacks and challenges from moving data management applications and tools into cloud systems. One important problem is how to adaptively modify the allocation of a workload within a cloud (Paton, N. W. et al. 2009). Other important problems are the configuration of virtual machines (Soror, A. et al. 2008), the automatic performance modeling of virtualized applications (Shivam, P., 2007) and self-tuning data management (Weikum, G. et al. 2002). The proposed framework is complementary to current research projects by working towards a holistic approach of moving data into the cloud and tackling the above problems: our goal is to profile the configuration of data services that can be offered by cloud providers and that are requested by consumers. It enables automation of data management performance, by offering data service composition and, finally, achieves pre-configured self-tuning data management via the realization of an open market of searching and matching offered and requested data services.

3 A FRAMEWORK FOR CLOUD DATA SERVICES

Problem Setting. The cloud computing paradigm is nowadays the answer for the ‘easy’ management of data, meaning the efficient processing and querying, especially of vast amounts of data, usually referred to as ‘big data’, but also various other data of many types, which are owned by small and medium enterprises and persons. Beyond the variety of data

[1] http://aws.amazon.com/simpledb/
[2] http://aws.amazon.com/rds/
[3] http://www.windowsazure.com/
[4] https://developers.google.com/cloud-sql/docs/introduction
[5] http://research.yahoo.com/project/212
[6] http://cassandra.apache.org/
[7] http://hbase.apache.org/
[8] http://srt-15.unine.ch

in terms of type and size, there is an enormous variety in the type of query workloads that people want to run on such data; in other words, the way that people want to access the, same even, data may vary tremendously in terms of duration and complexity. A *data service* in the cloud is the transparent management of data on a cloud infrastructure, and this may include data storage, maintenance and query execution.

There are three elements to combine in order to offer and demand data services, i.e. (a) infrastructure (b) data and (c) workload. These three elements need to be combined in an optimal way, in order to achieve the provision of optimal data services. Infrastructure, data and workload may either belong to the same person or organization or they may belong to different ones. In any case, the creation and the provision of a data service include instances of all three elements: a workload that is executed on some data using some part of the infrastructure. A data service may be requested by a user that owns the workload for execution or may be offered by the cloud that owns the infrastructure and/or the data. Data and workloads need services that comply with their management and execution characteristics, and cloud infrastructure can be employed to offer services that comply with its operational characteristics and availability.

We envision a virtual environment that constitutes an open market of cloud data services. In such an environment (Figure 1), cloud providers and their potential customers can advertise the offered and requested data services in a free manner and make contracts for service provisioning. Customers can be persons and organizations, and, even more, other cloud providers, or brokers of meta-services. The openness of the market aims at eliminating the current imbalance between consumers and producers of cloud data services, by allowing everybody to contribute and request services and enabling the near-real time negotiation of data service provision.

Motivation. Let us assume a computing situation with cloud infrastructure (IaaS) providers, data management (DBMS) providers and customers with data and/or workloads. The DBMS providers are the customers of the IaaS providers and data/workload owners are the customers of the data management providers. An open market of data services would allow all three parties to communicate and advertise their needs and availability, enabling efficient such matches. The participants of the market need to know what they are seeking and what they can offer: On one hand, the DBMS providers need information

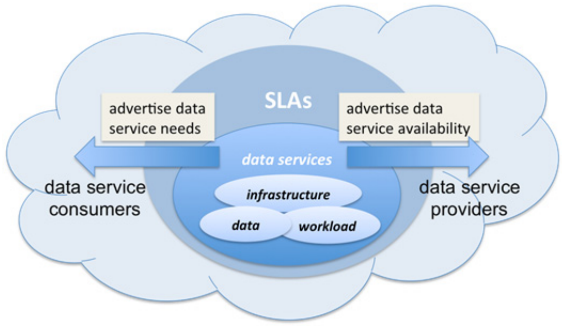

Figure 1: Efficient data service provisioning.

on the available choices for services from IaaS providers in order to form and build their data management services, and the data/workload owners need information on the available choices for data management services. On the other hand, the DBMS providers need to understand how the available infrastructure services can be used to build data management services, therefore, what data services can be built; moreover, data/workload owners need to understand the management and execution needs of their data and workload, respectively, therefore, what data services they are seeking. Also, the participating parties of the open market need to be able to advertise their offer or request for services in a coherent and comprehensive manner. Furthermore, they need to be able to search and match service needs and availability in a timely and efficient manner. Figure 2 shows graphically the described cloud environment.

Problem Definition. The provision of efficient cloud data services necessitates a successful match of offered and demanded services on infrastructure, data and workload. This requires means to (a) formulate, (b) express and (c) advertise offer and demand so that interesting parties can have the information needed for the selection of services.

a) Service Formulation: The basic requirement for the success of an open service market is for all involved parties to be aware of what are the services they can offer or demand, given some conditions. Specifically:

- Given some infrastructure and data, what are the workloads that can be served.
- Given some data and workload, what is the needed infrastructure.
- Given some workload and infrastructure, what should be the data to manage, i.e. the deployment of data and data structures.

b) Service Expression: The service matching requires the systematic and homogeneous expression of services from both ends of service provisioning,

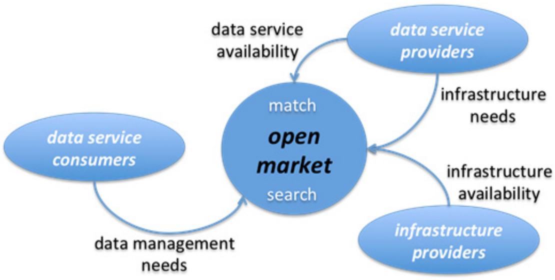

Figure 2: Open market of cloud data services.

i.e. the provider and the consumer, concerning all three elements, infrastructure, data and workload. Specifically, we need to know:

-The characteristics and metrics of possible services.

-The structural relations of services.

-A systematized construction of complex services.

c) Service Advertisement and Selection: Accurately formed and homogeneously expressed services can be advertised as the objects of search and match for service provision in the open market. The success of service match is based on an informative selection of services based on their provisioning characteristics. To achieve this, we need to know:

- How to efficiently and accurately match data service needs and availability.

- How to coordinate service advertisement in order to achieve near-real time service search and match.

Proposed Solution. The solution that we propose in order to achieve the creation and successful operation of an open market for cloud data services consists of an overall framework for the creation and management of data services, which includes: a technique for profiling data services; a methodology and methods for designing service offers and demands, and methods and algorithms for search and match of data services.

a) A Technique to Profile Data Services: Realistic data management needs and infrastructure capabilities can lead to the creation of a data service model. This should be based on the identification and standardization of functional characteristics and metrics for infrastructure, data and workload, as well as their relationships.

b) A Methodology and Methods to Design Data Services: The model enables the design of new services structured on service units. We need to:

- Define the service units as basic meaningful instances of the service model.

- Create a methodology for hierarchical service composition on top of service atoms and define composition directions and optimization objectives.

- Create specific methods for the hierarchical composition of examples of data services types.

c) Methods and Algorithms to Search and Match Data Services: The framework must be completed with the design of a service advertisement broker, a service matcher and a service scheduler.

- The service broker is the meeting point for producers and consumers of services, who are able to advertise their availability and need for services, respectively. Broker instances are able to communicate and form an overlay network for searching data services.

- The service matcher enables the bottom-up matching of offered and demanded services according to the respective instances of the data service model, employing a novel near-optimal matching algorithm for service profiles.

- The service scheduler takes as input the search offers and requests and schedule new, or renewals of, service matches in time frames in which these are valid. The scheduler must be based on novel methods that can monitor inputs of service advertisements. These methods will aim to change the scheduling time frames dynamically, in order to allow for optimal consumption of services through more accurate search and match.

The discussed framework aims at exploring the limits and capabilities of data service formulation and expression, and tailoring these to data service provisioning.

4 PROFILING DATA SERVICES

In order to formulate data services, the notion of a data service needs to become tangible. The first step is to study the data management needs and capabilities of cloud service consumers and providers, and to employ this study to propose a novel model for the description of a wide range of data services. Such a model can be employed to profile the data management needs and capabilities in order to form respective offered and demanded data services in a cloud environment. The goal is to produce a profiling technique that will allow for homogeneous service descriptions across data applications and platforms. Profiling the data services enables the creation of techniques for search and matching of services in an optimal manner.

Towards this direction we need to identify and study the key characteristics of data services. Such a study can give us a concrete idea of what is a realistic data service, what applications it can serve and in what way, what is the extent and the duration of the service, etc. The key characteristics of data services should be studied for the three elements of infrastructure, data and workload. Examples of basic characteristics that we expect to strongly determine the 'attitude' of services are:

- Concerning Infrastructure: CPU utilization, I/O operations, bandwidth, storage space, degree of shareness, size of virtual machines, etc.
- Concerning Data: replication degree, update rate, security constraints, data structures (views and indexes), partitioning degree and type, data statistics on selectivity estimation, etc.
- Concerning Workload: parallelization of execution, query complexity - translated into execution cost of query plans, data access skewness or similarity, etc.

The identification of the characteristics of data services nees to be accompanied with the definition of respective metrics. The service model should incorporate standard metrics such as response time, throughput and latency, as well as monetary cost. We also need to define new metrics, such as availability and privacy degree, and accompany each metric with an appropriate cost model, depending on the characteristic in hand.

5 DESIGNING DATA SERVICES

The profiling technique is the input to a novel methodology and methods for the design of cloud data services. The goal is to enable the creation of basic and generic data services, i.e. *service units*, and, furthermore, the hierarchical creation of complex data services on top of service units. We envision service units that comprise basic parts, qualitative or quantitative, of all three service elements, i.e. infrastructure, data and workload, profiled appropriately. For example, a service unit can be the combination of: a SPJ (i.e. select-project-join) query with one join, on an attribute with selectivity '10%' on a table of 10M tuples running on 1 CPU with no data transfer via network. Exploring realistic data services can lead to the proposal of example service units.

Furthermore, we need a methodology that facilitates users, meaning service providers and consumers, to create their own complex services along a range of optimization dimensions and ultimate service goals. Such a methodology can indicate paths of hierarchical service composition along such dimensions. We need to consider optimization dimensions respective to the metrics of service characteristics, such as response time and monetary cost. The methodology will take into account the durability or volatility of services, i.e. how persistent they are in time. Moreover, it will take into account any relationship between services. Such relationships can be either dependencies or incompatibilities between services, and may constrain their composition. Service relationships may be found either between instances of the same service element, or different ones. For example, it is incompatible to compose services with different data partitioning type, or a service that parallelizes workload execution depends on a service that updates data on all data replicas, respectively. Using the methodology, we will develop methods for hierarchical composition of services with specific optimization objectives and characteristics.

Since the service design is going to be employed in a dynamic cloud environment with frequent changes of available and requested data services, the goal of the composition methodology and methods will be to enable fast, i.e. near-real-time, composition based on data service patterns and re-usage of composition objects.

6 FINDING DATA SERVICES

To realize an open market of services we need to be able to search and match services in an accurate and efficient manner. Producers and consumers of services will be able to meet on brokers that implement the overall framework and advertise their availability and need for services via the respective instances of the data service model. They will be able to know their availability and need for services by using the profiling scheme and the service design methods. The search and match will target to produce valid combinations of the three service elements: infrastructure, data and queries. Each service to be searched or matched will have two parts, namely the *fixed* part and the *sought* part. The fixed part includes the conditions/constraints of request or availability, and the sought part of the service includes the actual request or availability to be matched. Naturally, we expect that, usually, the profiled infrastructure will be included in the sought part and workload will be included in the fixed part, for demanded services; we expect the opposite for offered services. Data can belong to either the fixed or the sought part, depending on the occasion. For example, if a service consumer needs to execute a workload on a specific database (e.g. workload and data belong to the same owner), then these data are included in the fixed part of the requested service. If a consumer seeks to execute a workload on some data that are out there, or she requests data services to deploy the data in the cloud, then these data may be included in the sought part of the service, since their deployment (i.e. storage, caching, built of data structures on top of it) in the cloud is flexible.

The overall framework can be implemented on service brokers of the open market. The framework instances can allow communication of the brokers by creating an overlay network (e.g. with centralized or Peer-to-Peer coordination), which can propagate local and received remote advertisements of services. We envision an open market that could take a further step down the road of searching and matching services, by enabling groups of users with semantically related needs for services to advertise their request or availability as a team and match services in a systematic manner.

While we expect search of services to be performed from top to bottom, meaning search performed based on the overall summary of complex hierarchically structured services, matching services in the framework should be enabled by algorithms that aims to match the services from bottom to top, meaning starting from the matching of atoms and moving to complex services built on top of them.

The goal is to enable dynamic and near-real-time search and match of services. Therefore, it is necessary to include methods that take as input search requests and service advertisements and schedule their matching in time frames in which the sought or advertised services are valid. Moreover, since the matched services may not have the same durability in time, the methods need to adapt the scheduling time frame in order to achieve optimal consumption of services through more accurate search and match. Such methods enable the maintenance of multiple dynamic service queues, the number and size of which depends on the variation of the durability of the incoming services.

7 CONCLUSIONS

The provisioning of data services is a new paradigm in data management that represents a very attractive solution for the management of business and large-scale data due to its low cost and high performance capabilities. The industry of cloud computing tries to catch up with fulfilling these data management needs but lacks the appropriate technology for the realization of cloud-hosted database systems, which is a critical component in the software stack of many cloud applications. The proposed framework aims to fill the gap between providers and consumers of data services that exists in today's cloud business, by not only solving the above two issues, but also by offering an all-inclusive solution for the offer of efficient and appropriate data services. Such a solution will enable the successful search and match of data services via the advertisement of service need and availability.

REFERENCES

Moore, R., Lopes, J., 1999. Paper templates. In *TEMPLATE'06, 1st International Conference on Template Production*. SCITEPRESS.

Smith, J., 1998. *The book*, The publishing company. London, 2nd edition.

Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., Zaharia, M., 2010. A view of cloud computing. *In CACM, vol. 53, pp. 50--58.*

DeCandia, G., Hastorun, D., Jampani, M., Kakulapati, G., Lakshman, A., Pilchin, A., Sivasubramanian, S., Vosshall, P., Vogels, W., 2007. Dynamo: Amazon's Highly Available Key-Value Store. *In Proc. of ACM SIGOPS, vol. 41, pp. 205--220.*

Chang, F., Dean, J., Sanjay, G., Hsieh, W. C., Wallach, D. A., Burrows, M., Chandra, T., Fikes, A., Gruber, R. E., 2006. Bigtable: A Distributed Storage System for Structured Data. *In Proc. of USENIX OSDI, pp. 5--15.*

Bernstein, P. A., Cseri, I., Dani, N., Ellis, N., Kalhan, A., Kakivaya, G., Lomet, D. B., Manne, R., Novik, L., Talius, T., 2011. Adapting Microsoft SQL Server for Cloud Computing. *In Proc. of ICDE, pp. 1255—1263.*

Kraska, T., Hentschel, M., Alonso, G., Kossmann, D. 2009. Consistency Rationing in the Cloud: Pay only when it matters. *In Proc. of VLDB, pp. 253--264.*

Abadi, D. J., 2009. Data management in the cloud: Limitations and opportunities. *In IEEE Bulletin on Data Eng., vol. 32, pp. 3—12.*

Aboulnaga, A., Salem, K., Soror, A. A., Minhas, U. F., 2009. Deploying database appliance in the cloud. *In IEEE Bulletin on Data Eng., vol. 32, pp. 13--20.*

Paton, N. W., Arago, M. A. T. de, Lee, K.,. Fernandes, A. A, Sakellariou, R., 2009. Optimizing Utility in Cloud Computing through Autonomic Workload Execution. *In IEEE Bulletin on Data Eng., vol. 32, pp. 51--58.*

Soror, A., Minhas, U. F., Aboulnaga, A., Salem, K., Kokosielis, P., Kamath, S., 2008. Automatic virtual machine configuration for database workloads. *In Proc. of ACM SIGMOD, pp. 953—966.*

Shivam, P., Demberel, A., Gunda, P., Irwin, D., Grit, L., Yumerefendi, A., Babu, S., Chase, J., 2007. Automated and On-Demand Provisioning of Virtual Machines for Database Applications. *In Proc. of ACM SIGMOD, pp.1079--1081.*

Weikum, G., Moenkeberg, A., Hasse, C., Zabback, P., 2002. Self-tuning Database Technology and Information Services: from Wishful Thinking to Viable Engineering. *In Proc. of VLDB, pp. 20—31.*