



A support for end user-Centric SLA administration of Cloud-Hosted Databases

A.Vasanthapriya¹, Mr. P. Matheswaran, M.E.,²

Department of Computer Science^{1,2}

ABSTRACT: Service level agreements for cloud services are not planned to flexible handle even comparatively uncomplicated act. In Exciting system, Cloud computing simplifies the time consuming processes of purchasing and software deployment. Service level agreements of cloud providers are not designed for supporting. The straight forward requirements and restrictions under which SLA of consumers applications need to be handled. We current approach for SLA based management of cloud hosted databases from the user perspective. We present an end-to-end framework for user centric SLA management of cloud hosted databases. Application defined policies for agreeable their individual SLA act requirements. Avoiding the cost of any SLA violation and calculating the economic cost of the allocated computing property. The end user applications with the mandatory flexibility for achieving their SLA supplies

KEY WORDS—Service Level Agreements (SLA), Cloud Databases, NoSQL Systems, Database-as-a-Service

I. INTRODUCTION

1.1 OVERVIEW

Cloud computing technology represents a new paradigm for the provisioning of computing infrastructure. This paradigm shifts the location of this infrastructure to the network to reduce the costs associated with the management of hardware and software resources. It represents the long-held dream of envisioning computing as a utility where the economy of scale principles help to effectively drive down the cost of computing infrastructure. Cloud computing simplifies the time-consuming processes of hardware provisioning, hardware purchasing and software deployment. Therefore, it promises a number of advantages for the deployment of data-intensive applications such as elasticity of resources, pay-per-use cost model, low time to market, and the perception of unlimited resources and infinite scalability. Hence, it becomes possible, at least theoretically, to achieve unlimited throughput by continuously adding computing resources (e.g. database servers) if the workload increases.

Cloud computing

Cloud computing technology represents a new paradigm for the provisioning of computing infrastructure. This paradigm shifts the location of this infrastructure to the network to reduce the costs associated with the management of hardware and software resources. It represents the long-held dream of envisioning computing as a utility [1] where the economy of scale principles help to effectively drive down the cost of computing infrastructure. Cloud computing simplifies the time-consuming processes of hardware provisioning, hardware purchasing and software deployment. Therefore, it promises a number of advantages for the deployment of data-intensive applications such as elasticity of resources, pay-per-use cost model, low time to market, and the perception of unlimited resources and infinite scalability. Hence, it becomes possible, at least theoretically, to achieve unlimited throughput by continuously adding computing resources database servers if the workload increases.

The store might want to know now what kind of things people buy together when they buy at the store. (If someone buys pasta, they usually also buy mushrooms for example.) That kind of information is in the data, and is useful, but was not the reason why the data was saved. This information is new and can be useful. It is a second use for the same data. Finding new information that can also be useful from data is called data mining.



International Journal of Innovative Research in Computer and Communication Engineering

(An ISO 3297: 2007 Certified Organization)

Vol.2, Special Issue 1, March 2014

Proceedings of International Conference On Global Innovations In Computing Technology (ICGICT'14)

Organized by

Department of CSE, JayShriram Group of Institutions, Tirupur, Tamilnadu, India on 6th & 7th March 2014

SLA

The mission of ITS is to provide high quality and reliable central Communications and Information Technology (C&IT) services that are cost-effective, based on best practice, and meet the requirements of College staff and students engaged in teaching & learning, research and administrative activities. The strategic context for the provision and development of central C&IT services is provided by the College's Communications & Information Technology, eLearning and Information Strategies, along with the ITS strategic and operating plans.

The Learning Technology team is primarily responsible for the promotion, development and support of the centrally provided learning management system (Blackboard via the BLE - Bloomsbury Learning Environment) and the provision of software application support on a range of core applications, particularly in areas that support the use of technology for teaching and learning. The Learning Technology support staff also plan, organise and run training workshops on a variety of general-purpose applications aimed at both student and staff users at all levels. Other responsibilities include writing user documentation, undertaking evaluation and procurement of new software packages and training materials, co-ordinating and promoting eLearning initiatives, providing support for staff development and undertaking help desk duties.

This theoretical approach is, necessarily, relational. An axiom of the social network approach to understanding social interaction is that social phenomena should be primarily conceived and investigated through the properties of relations between and within units, instead of the properties of these units themselves. Thus, one common criticism of social network theory is that individual agency is often ignored although this may not be the case in practice (see agent-based modelling). Precisely because many different types of relations, singular or in combination, form these network configurations, network analytics are useful to a broad range of research enterprises. In social science, these fields of study include, but are not limited to anthropology, biology, communication studies, economics, geography, information science, organizational studies, social psychology, sociology, and sociolinguistics.

Service Level

Service Level analysis or Service Levelling is the task of grouping a set of objects in such a way that objects in the same group called Service Level are more similar in some sense or another to each other than to those in other groups Service Levels. It is a main task of exploratory data mining, and a common technique for statistical data analysis used in many fields, including machine learning, pattern recognition, image analysis, information retrieval, and bioinformatics. Service Level analysis itself is not one specific algorithm, but the general task to be solved. It can be achieved by various algorithms that differ significantly in their notion of what constitutes a Service Level and how to efficiently find them. Popular notions of Service Levels include groups with small distances among the Service Level members, dense areas of the data space, intervals or particular statistical distributions.

Service Levelling can therefore be formulated as a multi-objective optimization problem. The appropriate Service Levelling algorithm and parameter settings (including values such as the distance function to use, a density threshold or the number of expected Service Levels) depend on the individual data set and intended use of the results. Service Level analysis as such is not an automatic task, but an iterative process of knowledge discovery or interactive multi-objective optimization that involves trial and failure. It will often be necessary to modify data pre-processing and model parameters until the result achieves the desired properties. Besides the term Service Levelling, there are a number of terms with similar meanings, including automatic classification, numerical taxonomy and typological analysis. The subtle differences are often in the usage of the results: while in data mining, the resulting groups are the matter of interest, in automatic classification primarily their discriminative power is of interest. This often leads to misunderstandings between researchers coming from the fields of data mining and machine learning, since they use the same terms and often the same algorithms, but have different goals.

We want to release aggregate information about the data, without leaking individual information about participants. Aggregate info: Number of A students in a school district. Individual info: If a particular student is an A

International Journal of Innovative Research in Computer and Communication Engineering

(An ISO 3297: 2007 Certified Organization)

Vol.2, Special Issue 1, March 2014

Proceedings of International Conference On Global Innovations In Computing Technology (ICGICT'14)

Organized by

Department of CSE, JayShriram Group of Institutions, Tirupur, Tamilnadu, India on 6th & 7th March 2014

student. We want to release aggregate information about the data, without leaking individual information about participants. Aggregate info: Number of A students in a school district. Individual info: If a particular student is an A student. Problem: Exact aggregate info may leak individual info. Number of A students in district, and Number of A students in district not named Frank McCrery.

We want to release aggregate information about the data, without leaking individual information about participants. Aggregate info: Number of A students in a school district. Individual info: If a particular student is an A student. Problem: Exact aggregate info may leak individual info. Number of A students in district, and Number of A students in district not named Frank McCrery. Goal: Method to protect individual info, release aggregate info. While cloud service providers charge cloud consumers for renting computing resources to deploy their applications, cloud consumers may charge their end users for processing their workloads or may process the user requests for free (cloud-hosted business application). In both cases, the cloud consumers need to guarantee their users' SLA. Penalties are applied in the case of SaaS and reputation loss is incurred in the case of cloud-hosted business applications. For example, Amazon found that every 100ms of latency costs them 1% in sales and Google found that an extra 500ms in search page generation time dropped traffic. In addition, large enterprise web applications (e.g., eBay and Facebook) need to provide high assurances in terms of SLA metrics such as response times and service availability to their users. Without such assurances, service providers of these applications stand to lose their userbase, and hence their revenues.

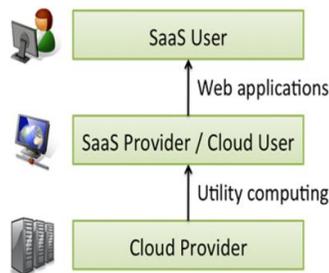


fig 1.2.1- Architecture of Cloud computing

1.3 Existing System

Cloud computing simplifies the time consuming processes of purchasing and software deployment. Service level agreements of cloud providers are not designed for supporting. The straight forward requirements and restrictions under which SLA of consumers applications need to be handled. SLA handling for their cloud hosted databases along with the limited of SLA frameworks. Service level agreements for cloud services are not designed to flexibly handle even relatively straightforward performance. Service level agreements for cloud services are not designed to flexibly handle .Most providers guarantee only the availability but not the performance of their services

1.4 Proposed System

End-to-end framework for consumer-centric SLA management of cloud-hosted databases. The framework facilitates adaptive and dynamic provisioning of the database tier of the software applications. Application defined policies for satisfying their own SLA performance requirements, avoiding the cost of any SLA violation. The framework continuously monitors the application-defined SLA. Automatically triggers the execution of necessary corrective actions when required. Cloud-based data management systems are to facilitate the job of implementing every application. A distributed, scalable and widely accessible service on the Web .The framework continuously monitors the application-defined SLA and automatically triggers the execution of necessary corrective actions.

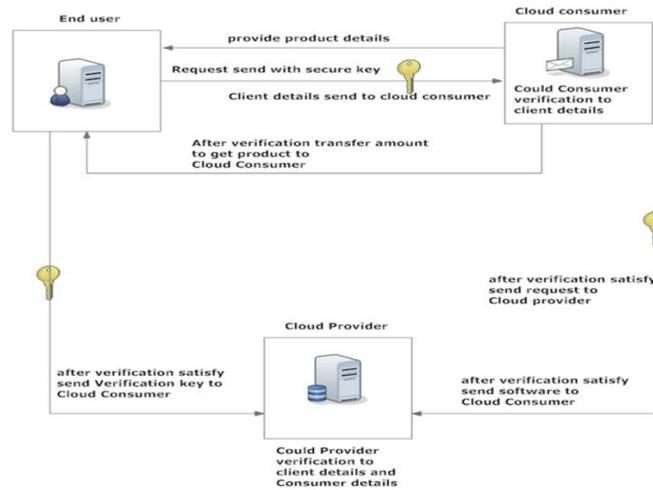


Fig 1.2.2 sla-system architecture

II. DATABASE REPLICATION IN THE CLOUD

2.1 THEOREM USED

CAP THEOREM

The CAP theorem shows that a shared-data system can, only choose at most two out of three properties: consistency (all records are the same in all replicas), Availability (all replicas can accept updates or inserts), and tolerance to Partitions (the system still functions when distributed replicas cannot talk to each other). It is highly important for cloud-based applications to be always available and accept update requests of data and at the same time cannot block the updates even while they read the same data for scalability reasons. Therefore, when data is replicated over a wide area, this essentially leaves just consistency and availability for a system to choose between. The sake of simplicity of achieving the consistency goal among the database replicas and reducing the effect of network communication latency, we employ the ROWA (read-once write-all) protocol on the Master copy. However, our framework can be easily extended to support the multi-master replication strategy as well.

2.2 PERFORMANCE EVALUATION

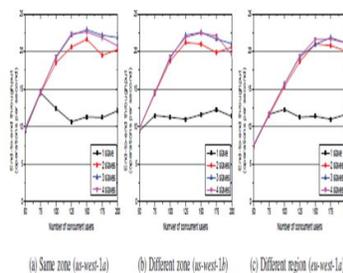


Fig. 4: End-to-end throughput of the workload with the read/write ratio 50/50.



International Journal of Innovative Research in Computer and Communication Engineering

(An ISO 3297: 2007 Certified Organization)

Vol.2, Special Issue 1, March 2014

Proceedings of International Conference On Global Innovations In Computing Technology (ICGICT'14)

Organized by

Department of CSE, JayShriram Group of Institutions, Tirupur, Tamilnadu, India on 6th & 7th March 2014

III. IMPLEMENTATION DETAILS

This project consists of modules that can be implemented by .Net and the modules are,

- 1)Service provider
- 2)End user Access and Product selection
- 3)Cloud Consumer Verification.
- 4)Cloud provider Verification
- 5)Download and analysis process

Service provider

Cloud service provider has been uploading all new products. In this upload product stored in cloud storage. This upload process any change only service provider .In this cloud storage secure to save for all product. If upload one product calculated all the details on that product. This work process for service provider work.

End user Access and Product selection

End user access process validate for all the details after that access process for next page to call. If end user create that time create account number generate if end user selected bank to automatically generate account number. After that enter your valid name and password to access process. And select software product selection process. If select software process and go to next page to get full details of product information .This is process for end user and product selection work.

Cloud Consumer Verification

Cloud consumer Verification process for verifying all information that end user. The cloud consumer gets secure information key to end user. That key to verify bank details and that user authorized bank person or to validate work for cloud consumer process. After validation process user select product amount is here or not check the end user bank details. Then all the details satisfy cloud consumer send secure key to Cloud provider .This work process for cloud consumer.

Cloud provider Verification

Cloud provider verification for cloud consumer details and End use details verification. Cloud consumer details verify and satisfy to cloud provider after that send verification key to end user. After that Software product send to cloud consumer. In this amount variation if end user to cloud consumer different amount fixed for owner. This work process for Cloud consumer verification process.

Download and analysis process

End user to download process for after key verification .If key receive to cloud provider to end user .After that verify key amount transfer to cloud consumer .Then download the end user select product. Then amount transfer cloud consumer to cloud provider .In this process to satisfy owner site and user side.finally analysis product and system performance.



International Journal of Innovative Research in Computer and Communication Engineering

(An ISO 3297: 2007 Certified Organization)

Vol.2, Special Issue 1, March 2014

Proceedings of International Conference On Global Innovations In Computing Technology (ICGICT'14)

Organized by

Department of CSE, JayShriram Group of Institutions, Tirupur, Tamilnadu, India on 6th & 7th March 2014

CONCLUSIONS

The design and implementation details of an end-to-end framework that facilitates adaptive and dynamic provisioning of the database tier of the software applications based on consumer-centric policies for satisfying their own SLA performance requirements, avoiding the cost of any SLA violation and controlling the monetary cost of the allocated computing resources. The framework provides the consumer applications with declarative and flexible mechanism for defining their specific requirements for fine-granular SLA metrics at the application level. The framework is database platform-agnostic, uses virtualization-based database replication mechanisms and requires zero source code changes of the cloud-hosted software applications.

REFERENCES

- [1] M. Armbrust and al., "Above the Clouds: A Berkeley View of Cloud Computing," University of California, Berkeley, Tech. Rep. UCB/EECS-2009-28, 2009.
- [2] J. Schad and al., "Runtime Measurements in the Cloud: Observing, Analyzing, and Reducing Variance," PVLDB, vol. 3, no. 1, 2010.
- [3] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears, "Benchmarking cloud serving systems with YCSB," in SoCC, 2010.
- [4] B. Suleiman, S. Sakr, R. Jeffrey, and A. Liu, "On understanding the economics and elasticity challenges of deploying business applications on public cloud infrastructure," Internet Services and Applications vol. 3, no. 2, 2012.
- [5] S. Sakr, A. Liu, D. M. Batista, and M. Alomari, "A survey of large scale data management approaches in cloud environments," IEEE Communications Surveys and Tutorials, vol. 13, no. 3, 2011.
- [6] W. Vogels, "Eventually consistent," Queue, vol. 6, pp. 14–19, October 2008. [Online]. Available: <http://doi.acm.org/10.1145/1466443.1466448>
- [7] R. Cattell, "Scalable sql and nosql data stores," SIGMOD Record vol. 39, no. 4, 2010.
- [8] H. Wada, A. Fekete, L. Zhao, K. Lee, and A. Liu, "Data Consistency Properties and the Trade-offs in Commercial Cloud Storage: the Consumers' Perspective," in CIDR, 2011.
- [9] D. Bermbach and S. Tai, "Eventual consistency: How soon is eventual?" in MW4SOC, 2011.
- [10] D. J. Abadi, "Data management in the cloud: Limitations and opportunities," IEEE Data Eng. Bull., vol. 32, no. 1, 2009.
- [11] D. Agrawal and al., "Database Management as a Service: Challenges and Opportunities," in ICDE, 2009.
- [12] S. Sakr, L. Zhao, H. Wada, and A. Liu, "CloudDBAutoAdmin: Towards a Truly Elastic Cloud-Based Data Store," in ICWS, 2011.
- [13] A. A. Soror and al., "Automatic virtual machine configuration for database workloads," in SIGMOD Conference, 2008.
- [14] E. Cecchet, R. Singh, U. Sharma, and P. J. Shenoy, "Dolly: virtualization-driven database provisioning for the cloud," in VEE, 2011.
- [15] P. Bodík and al., "Characterizing, modeling, and generating workload spikes for stateful services," in SoCC, 2010.
- [16] D. Durkee, "Why cloud computing will never be free," Commun. ACM, vol. 53, no. 5, 2010.
- [17] T. Ristenpart and al., "Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds," in ACM CCS, 2009.
- [18] C. Plattner and G. Alonso, "Ganymed: Scalable Replication for Transactional Web Applications," in Middleware, 2004.
- [19] A. Elmore and al., "Zephyr: live migration in shared nothing databases for elastic cloud platforms," in SIGMOD, 2011.
- [20] Y. Wu and M. Zhao, "Performance modeling of virtual machine live migration," in IEEE CLOUD, 2011.