

PERFORMANCE EVALUATION OF NOSQL CLOUD DATABASE IN A TELECOM ENVIRONMENT

Rasmus Paivarinta

Ixonos Plc, P.O. Box 284, FI-00811 Helsinki, Finland

Yrjo Raivio

Department of Computer Science and Engineering, Aalto University, P.O. Box 15400, FI-00076 Aalto, Finland

Keywords: NoSQL, IaaS, Home location register, Performance, SLA.

Abstract: Although the cloud computing paradigm has emerged in several ICT areas, the telecommunication sector is still mainly using dedicated computer units that are located in operators' own premises. According to the general understanding, cloud technologies still cannot guarantee carrier grade service level. However, the situation is rapidly changing. First of all, the virtualization of computers eases the optimization of computing resources. Infrastructure as a Service (IaaS) offers a complete computation platform, where instances can be hosted locally, remotely or in a hybrid fashion. Secondly, NoSQL (Not only SQL) databases are widely used in the internet services, such as Amazon and Google, but they are not yet applied to telecom applications. This paper evaluates, whether cloud technologies can meet the carrier grade requirements. IaaS cloud computing platforms and HBase NoSQL database system are used for benchmarking. The main focus is on the performance measurements utilizing a well known home location register (HLR) benchmark tool. Initial measurements are made in private, public and hybrid clouds, while the main measurements are carried out in Amazon Elastic Compute Cloud (EC2). The discussion section evaluates and compares the results with other similar research. Finally, the conclusions and proposals for the next research steps are given.

1 INTRODUCTION

Telecommunications operators are used to running their embedded computer systems on proprietary platforms. Typically operators have not shared infrastructures either, but have purchased their own networks. However, this situation is slowly changing. The first step has been taken by the Mobile Virtual Network Operators (MVNO), who have outsourced some part, or even the whole network, to network vendors. MVNOs have also utilized shared radio access networks (RAN) to avoid high initial investment costs. Recently, due to saturated revenues, cost pressures on operability and introduction of flat network architectures, such as Long Term Evolution (LTE), also dominant operators have shown interest in network sharing initiatives.

Cloud computing offers a new perspective on mobile network optimization. Unlike the past,

mobile networks are based on commercial computers equipped with the Linux operating system. Parallel to this, CPU and data storage performance are still developing almost exponentially (Armbrust, et al., 2009). This paradigm shift will open novel opportunities for cloud technologies in the telecom sector (Gabrielsson, et al., 2010).

It is probable that mobile networks will not change from private and proprietary servers to public, generic purpose computers in the short term, but telecom networks definitely include areas where cloud options can have a role. Especially mobile application servers and backend support systems might suit cloud computing well. The main drivers for the successful introduction of cloud technologies imply a large variation in traffic patterns or massive data volumes. In addition, telecommunication networks are normally designed based on the peak load, meaning that during off-peak periods systems

have a lot of unused capacity. Cloud computing thus offers a natural technology for resource sharing.

However, there are still concerns whether cloud computing meets the carrier grade requirements (Murphy, 2010). Service level agreements (SLA), such as high availability (HA), latency and transactions per second, are strict in several telecom services. One of the most critical mobile network functionalities is the home location register (HLR). HLR is the core element of the mobile system, and for example, is responsible for subscriber authentication and roaming functionalities. HLR also incorporates a risk for single point of failure, resulting in very high SLA requirements.

The paper evaluates, whether cloud technologies can meet the strictest telecom SLA requirements. HLR is used as a use case, although it is clear that the HLR, being the crown jewel of the operator, will not be the first functionality that operators would outsource to the cloud. However, the HLR presents exact SLA requirements, and also benchmark data and tools are available from the existing systems. HLR behavior in the cloud is studied by using two cloud technologies. First of all, all computation is placed into the Infrastructure as a Service (IaaS). IaaS can be applied locally, remotely and in both ways, referring to private, public and hybrid clouds, respectively. Secondly, the HLR benchmark tool is implemented in the HBase cloud database, which is based on NoSQL technology.

The paper is organized as follows. Firstly, the background data for the applied cloud technologies, namely IaaS and NoSQL, and also the benchmark tool, are presented. This is followed by the measurement system description. The results are shown using various setups, but the main emphasis is on the public cloud environment. The main criteria are latency and transactions per second. Next, the results are discussed, critically reviewed, and also a short business comparison is given. Finally, the conclusions and the future research proposals are made.

2 BACKGROUND

2.1 IaaS

An IaaS provides the most natural approach for the research. The existing telecom network elements, using the Linux operating system, can be easily ported as such into the IaaS platforms. Compared to the Platform as a Service (PaaS) or Software as a Service (SaaS) alternatives, IaaS offers the best

flexibility for its users. Unlike IaaS, PaaS service providers, such as Google App Engine or Microsoft Azure, require that the software is tailored for the associated platform. On the other hand, SaaS provides a complete service that does not allow running your own code. In addition, IaaS supports a large selection of open source software solutions that are compatible with the commercial IaaS market leader, Amazon Elastic Compute Cloud (EC2) (Amazon, 2011). By selecting the IaaS approach, the users can avoid the vendor or system lock-in, a feature that is much appreciated by the operators.

From commercial, public IaaS cloud vendors EC2, being a market leader, was a natural choice. On the private cloud side, the selection process was a lot more difficult. There are several alternatives in open source software IaaS platforms. The most well known, EC2 compatible, projects are called Eucalyptus (Eucalyptus, 2011), OpenNebula (OpenNebula, 2011) and OpenStack (OpenStack, 2011). As one of the more mature projects, Eucalyptus was selected for the private cloud platform, but for future research, OpenNebula and OpenStack are worth closer consideration.

Interoperability and backward compatibility of the software are essential features. Amazon EC2 and Eucalyptus provide an attractive duopoly, where software can be ported with minor efforts from one entity to another. The good interoperability basically enables two different scenarios. First of all, companies may develop their product using their own cloud, and at once a stable phase has been achieved, the software can be commercialized using a public cloud. The second possibility is to utilize a hybrid model, where private and public clouds complement each other, enabling load balancing functionalities. This is a valid scenario also in telecom applications, where the traffic peaks are a common challenge.

2.1.1 Amazon EC2

Amazon EC2 also provides several features that are important for the research carried out. First of all, EC2 offers a large variation of Linux distributions and instance types from small instances up to high performance computing (HPC) clusters. Secondly, the basic SLA guarantee, 99.95 percent, can be increased by using several, parallel, availability zones. The use of parallel zones is complemented by elastic IP addresses and monitoring services that support the implementation of HA targets. Thirdly, the hybrid clouds are backed by Virtual Private Clouds (VPC), securing the connections from local

clouds to remote public clouds. Finally, EC2 pricing structure is very flexible and enables various business models. For research purposes, EC2 is an affordable choice, because the pricing is based on the active computing hours.

2.1.2 Eucalyptus

Eucalyptus uses a set of services running on various terminals on different or same networks to manage and coordinate the whole system. A Cloud Controller (CLC) acts as a command terminal, which defines the cloud identity and resources available to it. It is the main service, which needs to be running prior to all other services of the cloud to function. A Cluster Controller (CC) provides management services for a set of clusters, controlled by a Node Controller (NC). It manages a set of end nodes where Eucalyptus Machine Images (EMIs) can run and support any user application. A Walrus storage service is a near clone of the EC2 Simple Storage System (S3). It provides a similar interface for storage and can use the same tools as are available for EC2 to manage the storage. Figure 1 illustrates the Eucalyptus architecture. (Nurmi, Wolski and Grzegorzczuk, 2009)

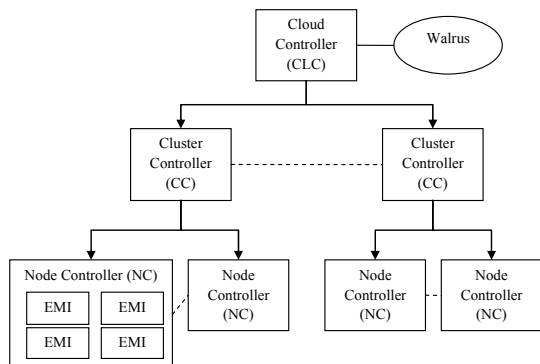


Figure 1: Eucalyptus architecture.

The architecture for setting up Eucalyptus can vary according to user needs. In a basic testing environment all four services can even reside on one machine and work together as a proper cloud deployment. However, there can be resource limitations depending on the hardware available. A better idea is to allocate one machine for CLC, CC and Walrus, and deploy NCs and EMIs on different computers.

2.1.3 Hybrid Cloud

Usually cloud services are implemented using public or private IaaS concepts. However, companies can

also choose a third deployment strategy, called a hybrid model. The basic idea is to utilize in parallel both public and private models. The selection can be done dynamically to match the current needs and to minimize the costs. The hybrid model sets a difficult selection challenge in the table. The technology should dynamically be able to provide load balancing between private and public clouds.

The hybrid model is not feasible for all applications and functions, but it looks attractive for services where the traffic load varies and the variations can be predicted well in advance. For example, a ticket sales service fulfils these criteria. Telecom networks also suffer from high traffic variations. Voice and text messaging services can become congested during exceptional events and on special dates. Figure 2 shows the text messaging volume trace on New Year's Eve (Zerfos, Meng, Wong and Samanta, 2006, p. 267). The data shows that during the midnight the peak load volume is tenfold compared to the average load. In a situation like this, hybrid clouds can offer an economic alternative for improving the end user experience during traffic peaks. Although the latency times can increase, a delayed service is a better option than no service at all.

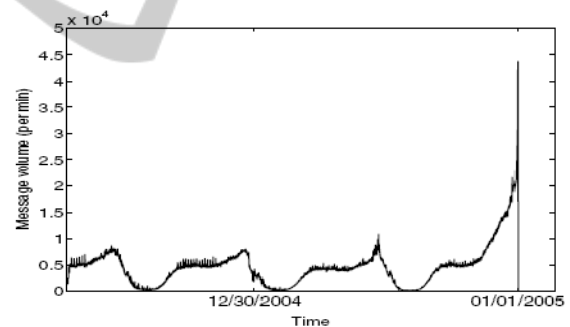


Figure 2: Text messaging traffic pattern at New Year (Zerfos, Meng, Wong and Samanta, 2006, p. 267).

2.2 NoSQL

Distributed databases have been at the forefront of cloud computing since the beginning, although the term NoSQL was invented much later. It is an umbrella term for a family of databases that typically do not implement the SQL interface, but are designed scale horizontally to support massive data. Originally the need to create a new kind of database stemmed from the data storage requirements of the first globally scale internet services. Soon, in addition to internal use at social media sites and internet companies, NoSQL solutions became available as services for all deve-

lopers.

The main differences between a NoSQL and a SQL, i.e. a Relational Database Management System (RDBMS), in a data model are provided interfaces, transaction guarantees and scalability. NoSQL differs fundamentally from the SQL databases that form the basis of telecom database systems. Generally RDBMS is optimal for online transaction processing (OLTP), and NoSQL for online analytics processing (OLAP) (Abadi, 2009). While a SQL database confirms ACID (atomicity, consistency, isolation, durability) requirements, NoSQL databases typically support BASE (Basically Available, Soft state, Eventually consistent) principles (Pritchett, 2008).

The modern history of the NoSQL movement as an effort to store web scale data can be seen to have begun in 2003 when Google published details on its Google File System (GFS) (Ghemawat, Gobioff and Leung, 2003). Later in 2006, the company published an article describing Bigtable (Chang, et al., 2006), a distributed storage system built on top of GFS. Imitating Google's efforts, the Apache Software Foundation (ASF) has developed open source clones, called the HBase (Apache, 2011) and Hadoop Distributed File System (HDFS) (Shvachko, Kuang, Radia and Chansler, 2010).

HBase and HDFS were chosen for a closer examination due to three reasons. First of all, HBase supports consistent transactions when updating a single row at a time. Secondly, it has a modular design and proven basis, thanks to underlying HDFS and ZooKeeper layers. Thirdly, HBase has active community and support from strong internet companies such as Yahoo. Yahoo has also developed a benchmark tool for cloud storages, including HBase (Cooper, et al., 2010).

2.3 TATP

The Telecommunication Application Transaction Processing (TATP) benchmark aims to measure the performance of a database under load which is typical in telecommunication applications. In particular, it is modelled after the type of queries that are processed in HLR on a GSM network. The benchmark tool is described in detail in the literature (Strandell, 2003; TATP, 2011). TATP encompasses seven different transactions of which three are reads and four are writes. The description gives probabilities at which each of the transactions is executed in the client. Broadly, 80 percent are reads and 20 percent are writes.

The database industry has been dominated by

RDBMSs for several decades, and it still is. Accordingly, TATP benchmark is heavily dependent on SQL, and does not provide functionality to test other kinds of database systems. However, we have taken action and implemented a comparable benchmark for HBase. The schema in TATP consists of four inter-relational tables. When modelling the schema for HBase, the tables were denormalised and finished off with just one table. Denormalisation is a popular approach when designing data models for NoSQL databases.

3 MEASUREMENT SETUP

3.1 Environment

The test environment simulates a real mobile network, where one HLR was loaded by one or several Mobile Switching Centers (MSC). The TATP benchmarking tool emulates the real signalling traffic between the MSCs and HLR. See Figure 3 for the model. The focus in the measurements was on the SLA, latency and transactions per second. HA measurements were beyond the scope of the research. All measurements were repeated a few times and the diagrams shown are based on average results.

It is noteworthy that the telecom level HA requirement, 99.999 percent, can be achieved by using independent IaaS clusters. For example, utilization of two different Amazon EC2 zones, both with an HA value 99.95 percent, yields together an HA value 99.9999 percent. Similar results can be achieved by using hybrid models.

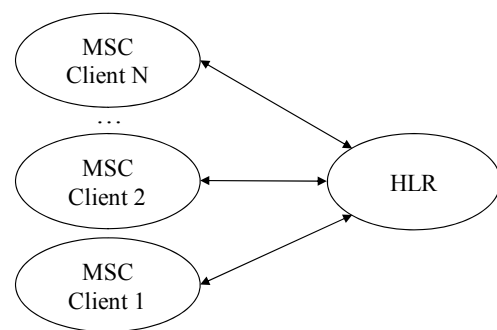


Figure 3: HLR test environment.

3.2 TATP Transactions

The TATP benchmark database representing HLR includes four different tables called Subscriber, Access Info, Special Facility and Call Forwarding.

Subscriber table includes the basic customer data, while Access Info table describes the access type. Special Facility table defines the services subscribers are entitled to, and finally Call Forwarding table reveals forwarding rules. MSC Clients use seven different transactions that either read or write data. The transaction distribution is known from real networks. Table 1 summarizes the transactions, their types, distribution and effected tables (Strandell, 2003).

Table 1: TATP transactions (Strandell, 2003).

Transaction name	Type	%	Tables
Get-Subscriber-Data	Read	35	Subscriber
Get-New-Destination	Read	10	Special Facility Call Forwarding
Get-Access-Data	Read	35	Access Info
Update-Subscriber-Data	Write	2	Subscriber Special Facility
Update-Location	Write	14	Subscriber
Insert-Call-Forwarding	Write	2	Call Forwarding
Delete-Call-Forwarding	Write	2	Call Forwarding

3.3 Initial Setup

The rewritten version of TATP was tested in several small HBase clusters. The focus was in transactions per second capability. Running the benchmark is interesting, especially because the results can be compared to existing reports on SQL database performance. One such article (Gupta, 2006) reports a throughput of approximately 5500 transactions per second. The performance level was achieved for 200 000 subscribers using carrier-grade hardware from the year 2006 and an in-memory database.

However, comparing measurement results obtained from different benchmarks testing different databases running on top of different infrastructures head-to-head, is not particularly meaningful. Therefore we use the results in the white papers only to set up a base line so that we know, whether the first results of running HBase in a HLR setting are on the same scale with recent commercial HLR databases.

In the initial measurements the environment was the following. The HBase version 0.20.6 and Hadoop 0.20.2 were run on a multitude of test setups, which all were considerably smaller than what HBase is designed for. Amazon Small EC2 had 1.7 GB memory on a Ubuntu Lucid 10.04 32 bit server. Eucalyptus had also 1.7 GB memory on top of a Ubuntu Lucid 64 bit desktop. Local

communication was based on a 100 Mbit/s LAN, and the PCs were equipped with Intel Core 2 Duo processors and 8 GB memory.

All setups consisted of a four virtual machine (VM) instance cluster. One instance was a dedicated master running the Hadoop Master Server and HDFS NameNode, two instances were running the HBase Region Server and HDFS DataNode processes, and the fourth machine was running a single HLR benchmark process and collecting the results. Instance deployment on Eucalyptus is presented in Figure 4. EC2 setup is similar.

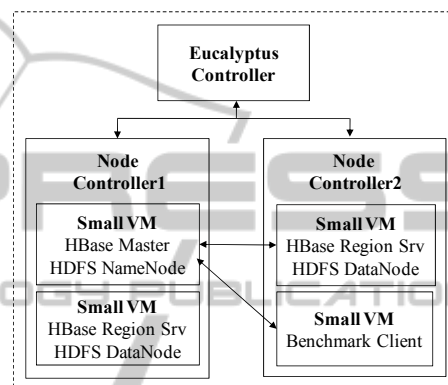


Figure 4: Test environment with Eucalyptus.

In the hybrid setup the HBase Region Servers and HDFS DataNodes were split into both EC2 and Eucalyptus, while the HBase Master, HDFS NameNode and benchmark client were running on a local Dell Optiplex 960 desktop. A noteworthy result itself is that we were able to run HBase and HDFS with default settings on Amazon Small EC2 instances without problems. Figure 5 presents the hybrid cloud setup.

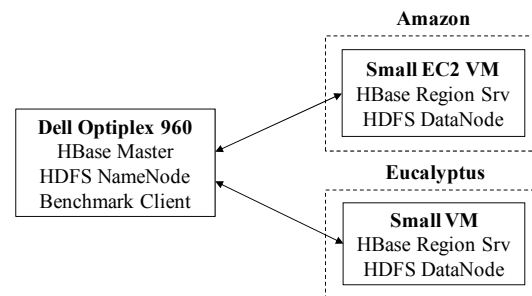


Figure 5: Hybrid cloud setup.

3.4 Final Setup

The final measurement setup was decided to be based on Amazon EC2 only. It was already beforehand clear that a hybrid IaaS architecture is

not optimal for a centralized database system. Furthermore, during the research process it was found that Eucalyptus is not the best platform for a hybrid cloud. The main reason is the lack of necessary management tools for remote instances. In contrast, OpenNebula and OpenStack support these functionalities, and for that reason those IaaS alternatives should be researched more in the hybrid context. On the other hand, private cloud measurement results were mainly limited only by the local hardware applied, having a small difference to the usual HLR environment.

In the later experiments, we investigated the effect of load, replication, database size and node failure on performance by running HBase on a cluster of six Large EC2 instances as shown in Figure 6. As the characteristics of EC2 instance types are sometimes modified, it is purposeful to specify here that the large instances used in the experimentation were virtual machines with 7.5 GB of memory and two virtual cores with two EC2 compute units, each running a 64-bit Ubuntu Server version 10.04. In addition to one master and four slave nodes, one large instance was hosting the benchmark clients.

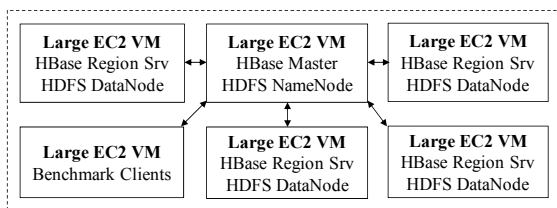


Figure 6: Final setup with Amazon Large EC2s.

4 MEASUREMENT RESULTS

4.1 Initial Measurements

The initial test results are shown in Figure 7. All transactions per second measurements were made with a database size of 200 000 subscribers. As described above, the comparison of different IaaS platforms with each other is not useful, but on the other hand, the results can be used for getting the big picture of the system. Compared to the four year old carrier grade numbers, the results were encouraging. The Amazon Large EC2 cluster achieved roughly 15 percent performance of the carrier grade system. Even the local, Eucalyptus based IaaS cluster managed to produce reasonable results. Plain workstation and legacy cluster measurements were made to gather experiences of running the bench-

mark setup in different environments.

The performance of a hybrid setup, consisting of Small Eucalyptus and EC2 instances, was better than with a Small EC2. The results prove that a hybrid IaaS can be made, and that the throughput is roughly the average of the building blocks. The first experiments also revealed that the bottleneck in the measurements was the single benchmark client. In the main measurements this bottleneck was removed by using several parallel clients. In the real networks one HLR is connected to several MSCs, too.

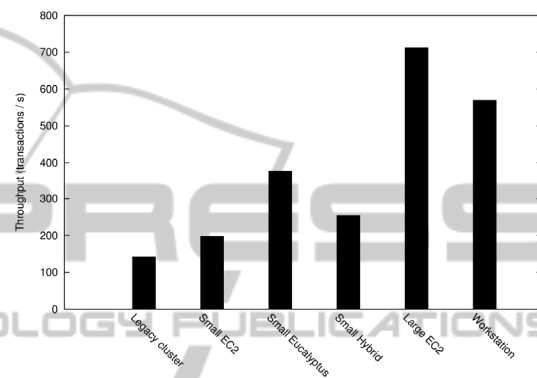


Figure 7: Throughput with various systems.

4.2 Main Measurements

The experiments analyzed here and presented in Figures 8 and 9 illustrate a load curve typical for I/O heavy systems. The throughput improves to a certain limit when client processes, i.e. load increases, are added, but after the limit latency grows dramatically. The point of maximal curvature is known as the knee. Our benchmark collects the latency distribution for each of the seven transaction types of TATP separately. Therefore, the response time values shown in the figures are the 95th percentile values perceived by the worst performing client process from the heaviest transaction type.

Replication is a standard way of achieving durability of data in NoSQL databases. Figure 8 shows the results of the experiment where the goal was to assess the effect of replication on performance. The table was populated with 200 000 subscribers. First of all, the throughput results show that 16 client processes are close to the knee, e.g. a point where the results turn worse. Secondly, we notice that the replication factor does not have a major impact on the throughput. Also the response time holds almost steady independent of the replication factor. We assume that the performance penalty of replication is virtually nonexistent, because even if writes become heavier, reads are

scaled across the replicas, which balances the results in a read heavy benchmark.

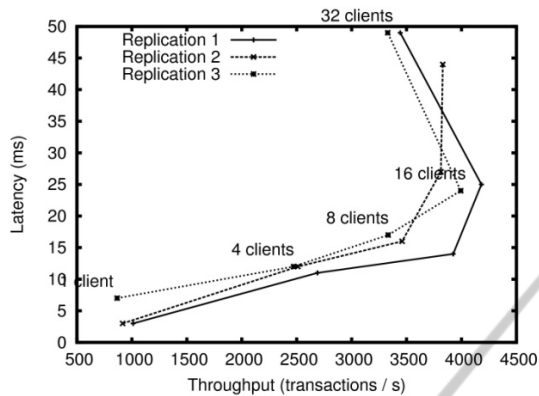


Figure 8: Impact of replication factor and number of clients to throughput and latency.

Figure 9 presents the effect of the amount of subscribers in the database on performance. The results were gathered when the replication factor was set to three. As expected, the performance gradually decreases as the table size grows. Looking at the results from 16 concurrent client processes, increasing the table size from one to five million subscribers decreases the throughput 32 percent and lengthens the response time 36 percent.

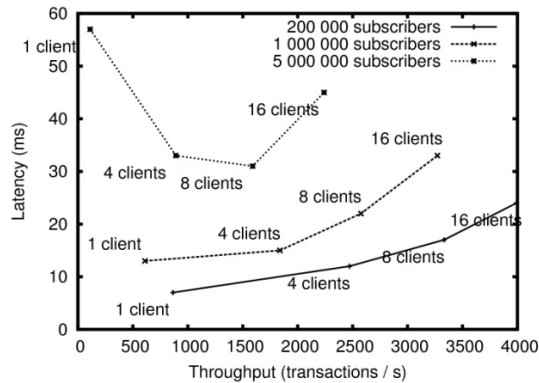


Figure 9: Impact of database size and number of clients to throughput and latency.

To verify that the HDFS replication gives protection from node failures, we studied the effect of killing one slave node in the middle of a benchmark run. The measurement was done with a database size of 1 million subscribers and replication factor two. Figure 10 shows the effect of one failing node 10 seconds after the launch of the run as perceived by four client processes. The throughput values are gathered once a second for each client and the results are stacked in the presentation. In this

sample the distributed database quickly recovers from the failure and continues serving clients within two seconds. The perceived recovery time in the experiment would be too much for real-time telecommunications applications, but it could be improved by tuning the parameters related to timeout mechanisms.

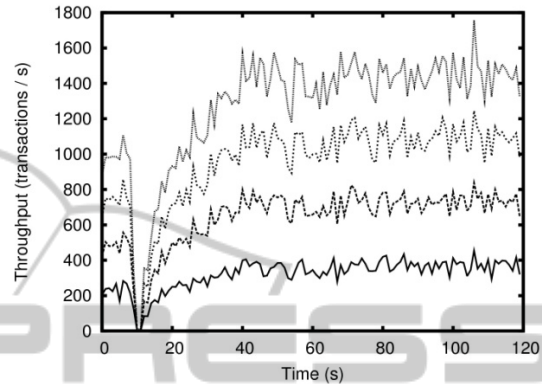


Figure 10: Recovery time from a node failure.

4.3 Measurements vs. Requirements

In order to give an idea of the load generated by the modified version of TATP, the performance testing tool bundled with HBase was also run on the test setup of six large EC2 instances. We run the performance test on the master node using 16 client threads and disabled MapReduce for it. By default, the test populates a table with one million rows of 1 kB each. In our experiment the random Read test took 1492 seconds, which leads to a throughput of 5.36 Mbit/s per client thread, and an aggregated throughput of 85.8 Mbit/s.

The 3GPP has defined the HLR performance requirements in their general specification (3GPP, 2009). According to that each subscriber produces on average 1.8 mobility related and 0.4 call handling related transactions per hour. Together this yields 2.2 transactions per hour per subscriber. With this information and the total number of subscribers, a requirement curve for transactions per second can be defined. Pulling together the requirements for HLR performance and the measurements from the HBase benchmark leads us to the conclusion that up to 4 million subscribers could in theory be supported by six large EC2 instances. Measurement results and comparable requirements are shown in Figure 11.

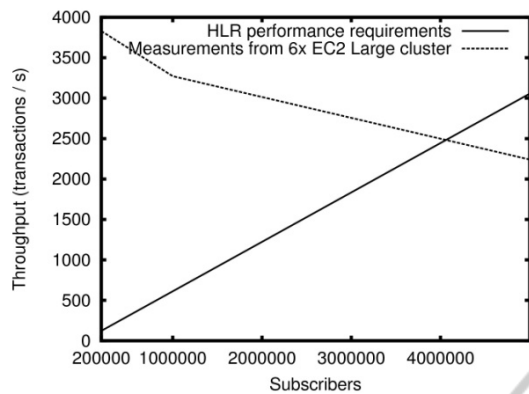


Figure 11: Number of subscribers and throughput vs. requirements.

In the experiment each subscriber added 3.7 kB to the database size leading to a total of 18.5 GB for 5 million subscribers. This is still in the area that can be handled in main memory, and therefore existing HLR solutions can support such deployments using an in-memory database. Similarly to most NoSQL databases, HBase does not support transactions, which span multiple rows, but on the other hand HBase guarantees that a single row remains consistent at all times. In an HLR all transactions read or update a single subscriber, and therefore the database was modelled so that all data related to a single subscriber is on one row.

5 DISCUSSION

Although the HLR is not a primary candidate for the cloud, the results give evidence that some other mobile network elements could be placed there. Application servers, such as SMS Center (SMSC), IP Multimedia Subsystem (IMS) and Service Delivery Platform (SDP), are examples of those. Backend processes can provide an even better solution area (Hajjat, et al., 2010). Billing, customer care and maintenance systems create a lot of data that could be computed by cloud infrastructures. A general purpose cloud can also be provided by mobile network vendors, who might use their large customer base to benefit from the statistical multiplexing. The same approach can work with operators, who operate in several countries and continents. It can be expected that in future mobile networks, such as Long Term Evolution (LTE), operators will compete and cooperate at the same time, leading to network sharing initiatives.

The HLR benchmarking measurements produced a lot of information about the cloud computing

opportunities in the telecom sector. The main lesson is that even the strict telecom SLA requirements can be achieved with both public and private clouds. The initial measurements also revealed that to match the existing RDBMS solutions, NoSQL databases have to fully utilize horizontal scalability. In addition, configuration parameters must be properly tuned, enterprise class infrastructure must be used and several client processes must be deployed.

IaaS, both private and public versions, operated according to expectations in the measurements. Due to the short history of private clouds, they are still developing. Amazon EC2, on the other hand, is already a mature product. The hybrid cloud was a side track in this research. It became evident that the hybrid cloud does not suit well to a centralized database system. In addition, the hybrid setup must be carefully designed to overcome configuration, management and load balancing challenges.

For certain applications a hybrid cloud can be an interesting option to optimize the dimensioning for peak loads (Moreno-Vozmediano, Montero and Llorente, 2009). However, the database solutions should be centrally located backed by 2N or N+1 redundancy algorithms. Database distribution will increase latency times and create unnecessary functional complications as well. For example, security and regulation challenges would become high.

The financial comparison of Amazon EC2, a private commodity server cluster and a carrier grade server would be interesting, but is not within the scope of this paper. The EC2 pricing structure is the most versatile, including also spot prices (Mattess, Vecchiola and Buyya, 2010). In a nutshell, EC2 prices are competitive with private clusters, especially if three-year term fixed prices and lower hour prices are utilized. The price scale of carrier grade services is large, which makes a reliable comparison almost impossible. It is also worth mentioning that, unlike in clouds (Greenberg, Hamilton, Maltz and Patel, 2009; Walker, Brisken and Romney, 2010), the weight of computing power and storage is marginal in the HLR price formula. However, in the application servers computing costs are becoming ever more dominant.

6 CONCLUSIONS

We have introduced research on how cloud computing performance meets the SLA requirements of mobile networks. The home location register (HLR) was chosen as an example for benchmarking

measurements. The HLR benchmark tool, originally developed for the SQL databases, was ported into the NoSQL, HBase specific environment. The software instances were deployed on private, public and hybrid Infrastructure as a Service (IaaS) platforms. The measurement results indicate that cloud technologies can achieve the mobile network latency and transactions per second requirements. Also telecom high availability (HA) targets can be met by using parallel computing zones. It is recommended that future studies should evaluate whether cloud technologies can be applied to mobile application servers and backend processes. Also the Long Term Evolution (LTE) will provide interesting research opportunities on network sharing between operators. Finally, the hybrid clouds deserve attention in managing traffic peaks.

ACKNOWLEDGEMENTS

The work is supported by Tekes (the Finnish Funding Agency for Technology and Innovation, www.tekes.fi) as a part of the Cloud Software Program (www.cloudsoftwareprogram.org) of Tivit (Strategic Centre for Science, Technology and Innovation in the Field of ICT, www.tivit.fi).

REFERENCES

- 3GPP, 2009. TS 43.005, version 9.0.0 release 9 - *Technical performance objectives*.
- Abadi, D. J., 2009. Data management in the cloud: limitations and opportunities. *IEEE Data Engineering Bulletin*, 32(1), March 2009.
- Amazon, 2011. *Amazon Elastic Compute Cloud (EC2)*. (online) Available at: [<http://aws.amazon.com/ec2/>](http://aws.amazon.com/ec2/) (Accessed: 24 February 2011).
- Apache, 2011. *The Apache Software Foundation – Hbase*. (online), Available at: [<http://hbase.apache.org/>](http://hbase.apache.org/) (Accessed: 24 February 2011).
- Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R. H., Konwinski, A., Lee, G., Patterson, D. A., Rabkin, A. and Stoica, I., 2009. Above the Clouds: A Berkeley view of cloud computing. *Technical Report No. UCB/EECS-2009-28*, Electrical Engineering and Computer Sciences, University of California at Berkeley, 10 February 2009.
- Chang, F., Dean, J., Ghemawat, S., Hsieh W. C., Wallach, D. A., Burrows, M., Chandra, T., Fikes, A. and Gruber, R.E., 2006. Bigtable: a distributed storage system for structured data. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI '06)*, Seattle, WA, USA 6-8 November 2006, USENIX Association.
- Cooper, B. F., Silberstein, A., Tam, E., Ramakrishnan, R. and Sears, R., 2010. Benchmarking cloud serving systems with YCSB. In *Proceedings of 1st ACM Symposium on Cloud Computing (SoCC '10)*, Indianapolis, Indiana, USA 10-11 June 2010.
- Eucalyptus, 2011. *The Open Source Cloud Platform*. (online) Available at: [<http://open.eucalyptus.com/>](http://open.eucalyptus.com/) (Accessed: 24 February 2011).
- Gabrielsson, J., Hubertsson, O., Más, I. and Skog, R., 2010. Cloud computing in telecommunications. *Ericsson Review*, vol. 1, pp. 29-33.
- Ghemawat, S., Gobioff, H. and Leung, S.-T., 2003. The Google file system. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP '03)*, New York, NY, USA 19-22 October 2003, ACM.
- Greenberg, A., Hamilton, J., Maltz, D. A. and Patel, P., 2009. The Cost of a Cloud: Research Problems in Data Center Networks. *ACM SIGCOMM Computer Communication Review*, Volume 39, Number 1, January 2009.
- Gupta, N., 2006. *Enabling High Performance HLR Solutions*. Netra Systems and Networking, Sun Microsystems, Inc., version 2, 17 April 2006.
- Hajjat, M., Sun, X., Sung, Y.-W.E., Maltz, D., Rao, S., Sripanidkulchai, K. and Tawarmalani, M., 2010. Cloudward bound: planning for beneficial migration of enterprise applications to the cloud. In *Proceedings of the 2010 ACM SIGCOMM*, New Delhi, India 30 August - 3 September 2010.
- Mattess, M., Vecchiola, C. and Buyya, R., 2010. Managing peak loads by leasing cloud infrastructure services from a spot market. In *Proceedings of the 12th IEEE International Conference on High Performance Computing and Communications (IEEE HPCC 2010)*, Melbourne, Australia 1-3 September 2010.
- Moreno-Vozmediano, R., Montero, R. S. and Llorente, I.M., 2009. Elastic management of cluster-based services in the cloud. In *Proceedings of the 1st workshop on Automated control for datacenters and clouds (ACDC '09)*, Barcelona, Spain 19 June 2009.
- Murphy, M., 2010. *Telco cloud*. (Presentation) Cloud Asia Singapore, 5 May 2010.
- Nurmi, D., Wolski, R. and Grzegorzczak, C., 2009. The Eucalyptus open-source cloud-computing system. In *Proceedings of the 9th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2009)*, Shanghai, China 18-21 May 2009.
- OpenNebula, 2011. *The Open Source Toolkit for Cloud Computing*. (online) Available at: [<http://www.opennebula.org/start>](http://www.opennebula.org/start) (Accessed: 24 February 2011).
- OpenStack, 2011. (online) Available at: [<http://openstack.org/>](http://openstack.org/) (Accessed: 24 February 2011).
- Pritchett, D., 2008. Base: an acid alternative. *ACM Queue*, vol 6, no. 3, pp. 48-55.
- Shvachko, K., Kuang, H., Radia, S. and Chansler, R., 2010. The Hadoop Distributed File System. In *Proceedings of IEEE 26th Symposium on Mass Storage*

- Systems and Technologies (MSST)*, Incline Village, NV, USA 3-7 May 2010, IEEE.
- Strandell, T., 2003. *Open source database systems: systems study, performance and scalability*. Master's thesis, University of Helsinki, Faculty of Science, Department of Computer Science and Nokia Research Center.
- TATP, 2011. *Telecom Application Transaction Processing Benchmark*. (online) Available at: <http://tatpbenchmark.sourceforge.net/> (Accessed: 24 February 2011).
- Walker, E., Brisken, W. and Romney, J., 2010. To lease or not to lease from storage clouds. *Computer*, IEEE, April 2010, pp. 44-50.
- Zerfos, P., Meng, X., Wong, S. H. Y. and Samanta, V., 2006. A study of the Short Message Service of a nationwide cellular network. In *Proceedings of Internet Measurement Conference 2006 (IMC '06)*, Rio de Janeiro, Brazil 25-27 October 2006.

