

YCSB+T: Bench-marking Web-Scale Transactional Databases

Varun Razdan¹, Rahul Jobanputra², Aman Modi³, Shruti Dumbare⁴

¹²³⁴Student, Dept. Of CE, Sinhgad Institute of Technology, Maharashtra, India

Abstract- *Yahoo Cloud Service Benchmark Client (YCSB). It can be used to benchmark new cloud database systems i.e. TPC-C and TPCE focus on emulating database applications to compare different DBMS implementations. It has ready adapters for different NoSQL Databases. YCSB allows benchmarking multiple systems and comparing them by creating "workloads". One can create and run on multiple systems on the same hardware configuration, and same workloads against each system. Many factors go into deciding which data store should be used for production applications, including basic features, data model, and the performance characteristics for a given type of workload. It's critical to have the ability to compare multiple data stores intelligently and objectively so that you can make sound architectural decisions. The Yahoo! Cloud Serving Benchmark (YCSB), an open source framework for evaluating and comparing the performance of multiple types of data-serving systems has long been the open standard for this purpose. In this paper, we designed a specific workload called Closed Economy Workload (CEW), which can run within the YCSB+T framework.*

In YCSB+T we develop new workload i.e. Closed Economy Workload (CEW) extended from workload from YCSB. As well as we concentrate on additional methods used to loads data or execute the workload on the database to validate its consistency. We observed that the number of transactions scales linearly up to 16 client threads. Our main motto is deal with data management access i.e. SELECT/UPDATE, with a large collection of items and operations that access and modify those items (get/put). We share our experience with using CEW to evaluate some NoSQL systems.

Key Words: Closed Economy Workload (CEW), TPC-C and TPC-E, AWS' S3, Scale-out, workload executor.

I. INTRODUCTION

There has been an explosion of new systems for data storage and management "in the cloud." Some systems are offered only as cloud services, either directly in the case of Amazon Simple DB and Microsoft Azure SQL Services, or as part of a programming environment like Google's App Engine or Yahoo!'s SQL. Still other systems are used only within. The large variety has made it difficult for developers to choose the appropriate system. The most obvious differences are between the various data models, such as the column-group oriented BigTable model used in Cassandra and HBase versus

the simple hash table model of Voldemort or the document model of CouchDB. However, the data models can be documented and compared qualitatively. Comparing the performance of various systems is a harder problem. Some systems have made the decision to optimize for writes by using on-disk structures that can be maintained using sequential I/O (as in the case of Cassandra and HBase), while others have optimized for random reads by using a more traditional buffer-pool architecture (as in the case of PNUTS). Furthermore, decisions about data partitioning and placement, replication, transactional consistency, and so on all have an impact on performance.

Understanding the performance implications of these decisions for a given type of application is challenging. Developers of various systems report performance numbers for the "sweet spot" workloads for their system, which may not match the workload of a target application.

II. PROBLEM CONTEXT

In context, we explain our work briefly to set of some common points as well as how extend YCSB to YCSB+T benchmark.

A. YCSB Benchmark

YCSB was developed at Yahoo! Labs to provide a framework and common set of workloads for evaluating the performance of different key-value stores. It has two parts:

- The YCSB Client, an extensible workload generator.
- The core workloads, a set of workload scenarios to be executed by the generator.

We care about query latency and overall system throughput. When take a closer look, however, the queries are very different. TPC-C contains several diverse types of queries meant to mimic a company warehouse environment. Some queries execute to the transactions over multiple tables; some are heavier in weight than others.

In contrast, the web applications we are benchmarking tend to run a huge number of extremely simple queries. Consider a table where each record holds a user's profile information. Every query touches only a single record, The YCSB "core

workload” accesses entire records at a time, and executes range queries over small portions of the table.

A.1.1) Workloads

YCSB includes a set of core workloads that define a basic benchmark for cloud systems.

The core workloads are a useful first step, and obtaining these benchmark numbers for a variety of different systems would allow you to understand the performance tradeoffs of different systems.

The core workloads consist of six different workloads:

Workload A: Update heavy workload

This workload has a mix of 50/50 reads and writes. An application example is a session store recording recent actions.

Workload B: Read mostly workload

This workload has a 95/5 reads/write mix. Application example: photo tagging; add a tag is an update, but most operations are to read tags.

Workload C: Read only

This workload is 100% read. Application example: user profile cache, where profiles are constructed elsewhere.

Workload D: Read latest workload

In this workload, new records are inserted, and the most recently inserted records are the most popular. Application example: user status updates; people want to read the latest.

Workload E: Short ranges

In this workload, short ranges of records are queried, instead of individual records. Application example: threaded conversations, where each scan is for the posts in a given thread (assumed to be clustered by thread id).

Workload F: Read-modify-write

In this workload, the client will read a record, modify it, and write back the changes. Application example: user database, where user records are read and modified by the user or to record user activity.

B. About NoSQL and Overview

The YCSB+T benchmark appears to be the best to date for measuring the scalability of SQL and NoSQL systems.

NoSQL databases are fast becoming the standard data platform for applications that make heavy use of telecommunication or Internet-enabled devices (i.e. browser-based, sensor-driven, or mobile) as a front end. While there are many NoSQL databases on the market, various industry trends suggest that the top three in use today are MongoDB, Apache Cassandra, and HBase.

Our experience with PNUTS tells us there are many design decisions to make when building one of these systems, and those decisions have a huge impact on how the system performs for different workloads (e.g., read-heavy workloads vs. write-heavy workloads), how it scales, how it handles failures, ease of operation and tuning, etc.

C. Other Recommendations and YCSB+T

YCSB+T is an important first step towards transactional benchmarking of scale-out data stores and provides a useful basis for quantifying the overhead introduced by wrapping CRUD operations in a transaction. However, the inclusion of the availability and replication tiers for transactional benchmarking which have already been proposed in the original YCSB.

NoSQL OLTP benchmarking is an active research topic and the boundaries of what is achievable in the field of distributed databases are being probed by both scientists and practitioners. Arguably, the most popular and most widely accepted OLTP benchmark for NoSQL databases is YCSB which facilitates measuring operational throughput and request latency for generic CRUD workload mixes.

In YCSB, read operations may read () a single row or scan () a range of consecutive rows and update operations may either insert () a new row or update () an existing one. Operations are issued one at a time per client thread and their distributions are based on parameters specified in the workload parameter files for a benchmark. The YCSB distribution includes five default workload files (called Workloads A, B, C, D and E) that generate specific read intensive, update-intensive and scan-intensive workloads. The current YCSB distribution provides DB client modules

with wrappers for HBase, Cassandra, MongoDB and Voldemort; YCSB++ adds a new client for Accumulo. Accumulo is the iterator framework that embeds user-programmed functionality into the different LSM-tree stages.

C.1.1) Architecture of YCSB+T

We implement YCSB+T which is extended by YCSB benchmark system. In YCSB+T we develop new workload i.e. Closed Economy Workload (CEW) extended from workload from YCSB. As well as we concentrate on additional methods used to loads data or execute the workload on the database to validate its consistency. We observed that the number of transactions scales linearly up to 16 client threads. Our main motto is deal with data management access i.e. SELECT/UPDATE, with a large collection of items and operations that access and modify those items (get/put).

The evaluation of YCSB+T demonstrates the usage for one particular system, but lacks a comparison of different transactional data stores. Benchmarking different scale-out transactional systems remains an important open issue in this field. YCSB+T furthermore does not detect transaction anomalies; it is limited to verifying state-based consistency constraints. The close economy workload is implement with the help of extend the workload class and some extra methods are implemented.

III. MATHEMATICAL EVOLUTIONS

Integrated mathematical steps necessary for implementation. The first expression to final all include input of our system with the help of mathematical parameter. In this section, we design mathematical expressions with the help of our existing and proposed system of our system.

A. Equations

The workload is most important parameter which is specifying by 'w'. Workload defines the data that will be loaded into the database during the loading phase, and the operations that will be executed against the data set during the transaction phase.

Typically, a workload is a combination of:

- Workload files.
- Parameter of workload file.

Because the properties of the dataset must be known during the loading phase (so that the proper kind of record can be constructed and inserted) and during the transaction phase (so that the correct record ids and fields can be referred to) a single set of properties is shared among both phases. Thus, the parameter file is used in both phases. The workload java class uses those properties to either insert records (loading phase) or execute transactions against those records (transaction phase).

For core workloads are 6. We can also implement new workload using our own parameter file with new values for the read/write mix, request distribution and extended our workload class.

$$W_c = \int_0^5 W_a + W_f \tag{1}$$

Where 'Wc' for core workload which is start from a up to f.

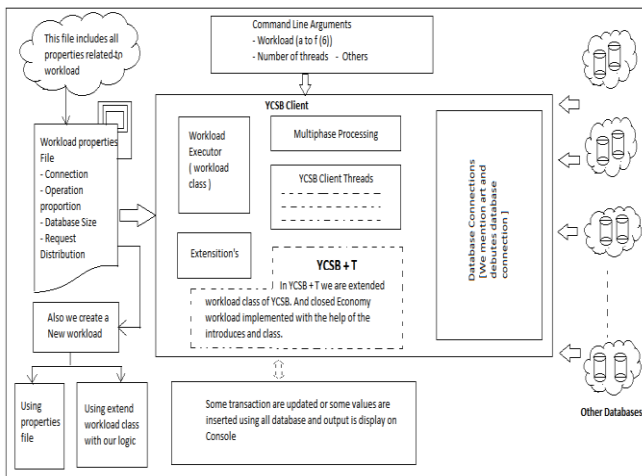


Fig -1: YCSB+T Architecture

D. Tiers and CEW

To the four tiers defined in the original YCSB contribution (performance, scalability, availability and replication) by introducing two new tiers transactional overhead and consistency. The Transactional overhead tier measures the latency of transactional operations (read, scan, insert, update, delete, read-and-modify) and transaction demarcation (start, abort, commit). To achieve this, a so-called Closed Economy Workload (CEW) is defined. It simulates bank account transactions in a closed system where money neither enters nor exits. This workload executes operations similar to YCSB, but wrapped in a single transactional context. For instance, doTransactionalReadModifyWrite reads two account records, transfers some money from one to the other and writes both records back. We have to create a new workload also.

$$W_{new} = \int_0^n W_0 + W_n \tag{2}$$

These workloads are choosing with specific appropriate runtime parameters i.e threads, target and status.

$$f(x) = \sum_0^n W_c \text{ or } \sum_0^n W_{new} \tag{3}$$

Where f(x) is the set of workloads which is come from our equation (1) and (2).

IV. UNITS

YCSB+T creates number of threads running on command line that can query the system which is under test. It can record latency in microseconds and calculate throughput in operations per second. We can measure replications, elasticity, availability, scalability and performance with YCSB+T benchmark which include extra transaction.

We are using same number of operation number and record number for every workload having range of 10000, 25000, 50000, 75000, 100000, 225000, 500000 and 1000000

(1 million) records. In each workload, there will be a loading stage in which insert operation takes place which is basically loading performance for each workload while it is easy to obtained running performance from every workload.

TABLE I
UNITS FOR MAGNETIC PROPERTIES

Symb ol	Explanation	Conversion from Second and Millisecond and Microsecond and Other
t	throughput	1 Ope/Sec → 10 ³ /Sec = 10 ⁶ Ms
l	latency	1 Msec → 10 ³ MiSec.
db	database	1 DB → Load Any One
W	workload	Choose basic or new.
RMW	Read modify	1 Ope/sec → 100103
COP	commit operation	1 Ope/Sec → 1000000

	with thread 16	
U	Update operation with thread 16	1 Ope/Sec → 200206
AOP	actual operation	1 Ope/Sec = 10 ⁵ /Sec
ASC	anomaly score	1 score/sec = 2.9E-5
OP	operations	1 Ope/Sec = 100000
CEW	Closed economy workload	Average Latency → 6134.37
μ _s	microseconds	μ _T → μ _s
MIN _{late}	maximum latency	1 → depend on thread
MIN _{late}	minimum latency	1 → depend on thread

Above tables shows all mathematical notions are used in development with their values.

V. CONCLUSION

The performance of the Proposed System and scalability of large-scale distributed NoSQL systems like Yahoo! PNUTS as well as traditional database management systems like MySQL. We implement YCSB+T which is extended by YCSB benchmark system. In YCSB+T we develop new workload i.e. Closed Economy Workload (CEW) extended from workload from YCSB. As well as we concentrate on additional methods used to loads data or execute the workload on the database to validate its consistency.

REFERENCES

[1] Wei Wei, Qi Yong. Information potential fields navigation in wireless Ad-Hoc sensor networks[J]. Sensors, 2011, 11(5): 4794-4807.

[2] Song, Houbing, and Maité Brandt-Pearce. "A 2-D discrete-time model of physical impairments in wavelength-division multiplexing systems." Journal of Lightwave Technology 30.5 (2012): 713-726.

[3] Ion Stoica, Robert Morris, David Liben-nowell, David R. Karger, M. Frans Kaashoek, Frank Dabek, and Hari Balakrishnan. Chord: a scalable peer-to-peer lookup

- protocol for internet applications. IEEE/ACM Transactions on Networking, 11:17–32, 2003.
- [4] Zhou Wei, Guillaume Pierre, and Chi-Hung Chi, “CloudTPS: Scalable Transactions for Web Applications in the Cloud,” *IEEE TRANSACTIONS ON SERVICES COMPUTING, VOL. 5, NO. 4, OCTOBER-DECEMBER 2012*
- [5] W. Vogels, “Data Access Patterns in the Amazon.com Technology Platform,” Proc. 33rd Int’l Conf. Very Large Databases (VLDB), 2007.
- [6] M. Brantner, D. Florescu, D. Graf, D. Kossmann, and T. Kraska, “Building a Database on S3,” Proc. SIGMOD Int’l Conf. Management of Data, pp. 251-264, 2008.
- [7] J. Gray, editor. The Benchmark Handbook for Database and Transaction Processing Systems. Morgan Kaufmann, 2003
- [8] “PL/SQL User’s Guide and Reference, Version 3.0,” Part 800-V1.0, Oracle Corp., 2004.
- [9] M. Stonebraker et al. C-store: a column-oriented DBMS. In VLDB, 2005.
- [10] Amazon SimpleDB. <http://aws.amazon.com/simplydb/>.
- [11] Google App Engine. <http://appengine.google.com>.
- [12] Yahoo! Query Language. <http://developer.yahoo.com/yql/>.
- [13] P. Shivam et al. Cutting corners: Workbench automation for server benchmarking. In Proc. USENIX Annual Technical Conference, 2008.
- [14] <http://www.ibm.com/developerworks/opensource/library/os-couchdb/index.html>. Exploring Couch DB.
- [15] Java development 2.0: MongoDB: A NoSQL datastore with (all the right) RDBMS moves. <http://www.ibm.com/developerworks/java/library/j-javadev2-12/>.
- [16] Bringing Big Data to the Enterprise. <http://www-01.ibm.com/software/data/bigdata/>.