

Active Cloud DB: A Database-Agnostic HTTP API to Key-Value Datastores

Chris Bunch Jonathan Kupferman Chandra Krintz
Computer Science Department
University of California, Santa Barbara
April 26, 2010 UCSB Tech Report 2010-07

Abstract

In this paper, we present a technique that connects web applications to cloud-based distributed datastore technologies that implement the Google App Engine cloud datastore API. We implement our approach as a Google App Engine (GAE) application that we employ to expose the GAE datastore API to developers – for use with any language and framework. We evaluate this application on both GAE and over AppScale, the open-source implementation of GAE that removes the programming and library restrictions of GAE and that enables GAE applications to execute on virtualized cluster resources, including Eucalyptus and EC2, without modification. As part of this work, we extend AppScale with simple caching support to improve the performance of datastore access and evaluate our technique with and without this support. We also make use of this support within multiple prototypes (e.g. Ruby/Rails, Python/Django) to show the ease-of-use and applicability of our contribution to other web development environments.

1 Introduction

Distributed key-value datastores have become popular in recent years due to their simplicity, ability to scale within web applications and services usage models, and their ability to grow and shrink in response to demand. As a result of their success in non-trivial and highly visible cloud systems for web services, e.g. BigTable [4] within Google, Dynamo [1] within Amazon, and Cassandra [3] within Facebook, a wide variety of open-source variations of distributed key-value stores have emerged and are gaining wide-spread use.

However, these datastores implement a wide variety of features that make them difficult for prospective users to compare. For example there are differences in query languages, topology (master/slave vs peer-to-peer), consistency policies, and end-user library interfaces. As a result, we and others have investigated a single frame-

work with which such systems can be compared. The two systems which enable this are YCSB [6] and the AppScale platform [5, 2]. YCSB is the Yahoo! Cloud Serving Benchmark. YCSB provides a new DB interface and a synthetic workload executor for exercising the DBs that the authors attach to the interface. The system measures the response time of primitive operations in a workload between a thread on one machine and the datastore on another. The authors have support for four datastores (without replication): HBase, Cassandra, PNUTS (Yahoo’s internal key-value store), and Sharded MySQL.

AppScale is an open-source implementation of the Google App Engine (GAE) cloud platform. It employs the GAE Datastore API as a unifying API through which any datastore can be “plugged in”. Once a datastore implementation is added to AppScale, it is deployed and configured automatically (there are command-line parameter settings for replication factor, cloud size, etc.) within the AppScale cloud deployment. Thus AppScale automates the configuration and deployment of these complex distributed systems and facilitates different databases to be compared. AppScale currently implements this key-value API using HBase, Hypertable, Cassandra, Voldemort, MongoDB, MemcacheDB, and MySQL Cluster. The system however, unlike YCSB, measures application end-to-end performance (the round trip time from the user/browser to the web server to the datastore, and back). Unfortunately, AppScale only supports applications written in the languages that GAE supports (currently Python and Java). YCSB does not support applications at all but instead spawns requests between the server and the datastore for the sole purpose of measuring datastore response time and throughput.

In this work, we address the problem of how to easily provide users of any programming language and framework with a database-agnostic interface to key-value datastores. To enable this, we (i) employ AppScale as a cloud platform to configure and deploy the needed datastore, and (ii) design and implement a GAE application

that runs over AppScale, called Active Cloud DB (in the spirit of Ruby’s ActiveRecord). Active Cloud DB is a cloud software component that exposes an HTTP API that is implemented via GAE datastore API.

As a result, Active Cloud DB provides application developers with automatic and straight-forward access to the scalable datastore back-ends that AppScale implements and to GAE BigTable. Active Cloud DB is easy to use and to integrate into language and framework front-ends and also provides the functionality to move data between distributed datastores without manual configuration or deployment and without knowledge about the datastore implementation.

We employ the API to implement support for disparate web-based language frameworks, including Ruby/Sinatra, Ruby/Rails, Python/Django, and Python/Web.py. We evaluate the performance of the API using Cassandra and MemcacheDB datastores to investigate the performance differences of the primitive datastore operations. Finally, we extend AppScale with a simple data caching scheme mechanism for Active Cloud DB – that also can be used by any GAE application that runs over AppScale, and evaluate its impact.

In the sections that follow, we overview AppScale, the implementation and evaluation framework we use to investigate the performance of our approach. We then describe Active Cloud DB and the caching support with which we extend AppScale. We then describe four proof of concepts we have developed that make use of Active Cloud DB, evaluate this web service, and conclude.

2 AppScale

AppScale is a robust, open source implementation of the Google App Engine APIs that executes over private virtualized cluster resources and cloud infrastructures including AWS and Eucalyptus [8]. Users can execute their existing App Engine applications over AppScale without modification. [5] describes the design and implementation of AppScale. We summarize the key components of AppScale that impact our description and implementation of Active Cloud DB.

AppScale is an extension of the non-scalable testing/debugging front-end that Google makes available as open-source, and to which Google refers as the Google App Engine (GAE) software development kit. This extension is called the AppServer in AppScale. AppScale integrates open-source datastore/database systems as well as new software components that facilitate configuration, one-button deployment, and distributed system support for scalable, multi-application, multi-user cloud operation.

The AppServer decouples the APIs from their non-scalable SDK implementations and replaces them dis-

tributed and scalable versions – for either Python or Java. That is, GAE applications can be written in either language. AppScale implements front-ends for both languages in this way.

GAE applications write to a key-value datastore using the following Google Datastore API:

- Put(k, v): Add key k and value v to table; creating a table if needed
- Get(k): Return value associated with key k
- Delete(k): Remove key k and its value
- Query(q): Perform query q using the Google query language (GQL) on a single table, returning a list of values
- Count(t): For a given query, returns the size of the list of values returned

GAE applications employ this API to save/restore data as part of the GAE application implementation. The AppServer (and GAE SDK) encodes each request using a Google Protocol Buffer [9]. Protocol Buffers facilitate fast encode/decode times and provide highly compact encoding. As a result, they are much more efficient for data serialization than other approaches. The AppServer sends/receives Protocol Buffers to/from a Protocol Buffer Server in AppScale over encrypted sockets. All details about and use of Protocol Buffers and the back-end datastore (in Google or in AppScale) is abstracted away from GAE applications using this API.

The AppScale Protocol Buffer Server implements all of the libraries necessary to interact with each of the datastores that are plugged into AppScale. Since this server interacts with all front-ends, it must be very fast and scalable so as not to impact end-to-end application performance. AppScale places a Protocol Buffer Server on each node to which datastore reads and writes can be sent (i.e. the datastore entry points) by applications. For master-slave datastores, the server runs on the master node. For peer-to-peer datastores, the server runs on all nodes.

AppScale currently implements seven popular open-source datastore/database technologies. They are Cassandra, HBase, Hypertable, MemcacheDB, MongoDB, Voldemort, and MySQL Cluster (with a key-value data layout). Each of these technologies vary in their maturity, eventual/strong consistency, performance, fault tolerance support, implementation and interface languages, topologies, and data layout, among other characteristics. [2] presents the details on each of these datastores.

Active Cloud DB and our caching support is datastore-agnostic. These extensions thus work for all of the AppScale datastores (current and future) and for GAE applications deployed to Google’s resources. In our evaluation section, we choose two representative datastores

(due to space constraints) to collect empirical results for: Cassandra and MemcacheDB. We provide a brief overview of each.

Cassandra. Developed and released as open source by Facebook in 2008, Cassandra is a hybrid between Google’s BigTable and Amazon’s Dynamo [3]. It incorporates the flexible column layout from the former and the peer-to-peer node topology from the latter. Its peer-to-peer layout allows the system to avoid having a single point of failure as well as be resilient to network partitions. Furthermore, read and write requests can be sent to any node in the system, allowing for greater throughput. However, this also results in data being “eventually-consistent.” Specifically, this means that a write to one node will take some time to propagate throughout the system and that application designers need to keep this in mind. Cassandra exposes a Thrift API through which applications can interact with in many popular programming languages.

MemcacheDB. Developed and released as open source by Steve Chu in 2007, memcachedb is a modification of the popular open source distributed caching service memcached. Building upon the memcached API, MemcacheDB adds replication and persistence using Berkeley DB [7] as a backing store. MemcacheDB provides a key-value datastore with a master-slave node layout. Reads can be directed to any node in the system, while writes can only be directed to the master node. This allows for strong data consistency but also multiple points of access for read requests. As MemcacheDB is a master-slave datastore, the master node in the system is the single point of failure. MemcacheDB can use any library available for memcached, allowing for native access via many programming languages.

3 Active Cloud DB

We next present Active Cloud DB, a software service that executes over AppScale to expose datastores to cloud clients. Clients are web-based applications and software implemented using language frameworks other than Google App Engine. Active Cloud DB exports to clients access to cloud-based distributed datastore technologies automatically and scalably.

Active Cloud DB is a software layer that we implement as a GAE application that executes over AppScale (and thus Google App Engine). Its implementation requires no modification to AppScale. To enable scalability and low response times, we also extend AppScale with datastore caching support to improve the performance of datastore access by applications. We first

overview Active Cloud DB and then describe the design and implementation of our caching scheme.

3.1 Implementation

Our goal with Active Cloud DB is to remove the limitation imposed by AppScale and Google App Engine that requires that applications be written in Python or Java to gain access to the functionality of the distributed datastore back-ends that these cloud fabrics implement. This is particularly important for Google App Engine, since not only must applications be written in these languages but they must employ the web frameworks specified by Google using a restricted set of “whitelisted” libraries. AppScale removes this constraint but currently still only supports Python and Java GAE applications.

Active Cloud DB is a Google App Engine application that runs over AppScale that exposes an HTTP API to the underlying datastore. Internally, Active Cloud DB allows for objects to be created with a given name (key) and one string within (its value). This can be trivially extended to allow for all data types, but we consider only string-type keys and values here.

Active Cloud DB exposes three URL routes: a `get` route, a `put` route, and a `delete` route, each of which map to their corresponding primitive operations. Applications perform an HTTP POST request with the necessary key or key/value pair via the AppServer. Requests can also be or the AppLoadBalancer provided by AppScale, which dynamically maps requests to AppServers.

3.2 Integration

To communicate with Active Cloud DB, developers implement the client side-interface. We have done this for four popular web frameworks, Rails and Sinatra for Ruby, and Django and web.py for Python. In Rails and Django, we remove the built-in database abstractions and in all add in a wrapper for communicating with Active Cloud DB. It abstracts all remote communication logic and error handling such that the application developer need not be aware that the database is located remotely. In all frameworks we use the default templating library for creating the presentation layer of the applications.

The four prototype applications with which we make use of Active Cloud DB implement a simple bookstore application (inspired by the application given in [10]). To access Active Cloud DB, the application makes HTTP POST requests with the key and value to put, or the key to get or delete. The bookstore application maintains a special key containing a list of all the books in the datastore, which is maintained whenever books are added or removed. When a book is added, we change this special key accordingly and then add an entry to the datastore with the book’s name (the key) and a summary of

the book's contents for users to view (the value). When users wish to view all books in the bookstore, we access the special key to get a list of all the books, and for each book (key), we return the corresponding book information (its value).

3.3 Caching Support

To reduce latency and improve throughput to/from the datastores in AppScale, we provide a transparent caching layer for data. To enable this, we leverage the Google Memcache API (similar to memcached). The API caches web service data and provides a least-recently-used replacement policy.

To efficiently cache data, we combine two caching strategies, write-through and generational. For basic operations (get, put, delete) we employ a write-through caching strategy. With this strategy all put operations are written to the datastore as well as to the cache. In doing so, subsequent requests will likely be served directly from cache, avoiding any datastore interaction.

Efficiently caching query operations is more complex than basic operations because query results can contain multiple data items. Furthermore, when a particular data item is updated one must expire all query results which contain that item so that stale data is not returned. In order to ensure this property we utilize a *generational* caching strategy. In essence, a generation value is maintained for the data. The generation value is included in the cache key for all query operations. Hence, by changing the generation value all prior cached results are implicitly expired as they will never be accessed again.

Specifically, the operations are:

- Get(k): Return value associated with key k . If the value was not in the cache, store it for future accesses.
- Put(k, v): Add key k and value v to table and the cache. Increment the generation value.
- Delete(k): Remove key k and its value from the table and the cache. Increment the generation value.
- Query(q): Perform query q using the Google Query Language (GQL) on a single table, returning a list of values. Store the result in the cache with the current generation number for future queries.
- Count(t): Acquire the query data via the new query technique, and return the size of the list of values found.

4 Evaluation

We next employ Active Cloud DB over AppScale to evaluate the performance characteristics of the various supported datastores. We begin by describing our methodology and the present our results.

4.1 Methodology

For our experiments, we measure the performance of the primitive operations performed in bulk and as part of an overall workload. In both scenarios, this is done over two back-end AppScale datastores, Cassandra v0.5.0 and Memcachedb v1.2.1-Beta. For the first set of experiments, we fill a table in each database with 1000 items and perform the get, put, query, and delete operations. To provide a baseline measurement we also perform a no-op operation which simply returns and performs no back-end processing. We invoke each operation 1000 times (100 times for query since it takes significantly longer than the others). A query retrieves all 1000 items from the datastore.

For each experiment, we access Active Cloud DB using a machine on the same network. Our measurements are of round-trip time to/from the AppServer as well as all database activity. For each experiment, there are nine concurrent threads that each perform all of the operations and record the times for each. We consider multiple static configurations of the AppScale cloud that consists of 4, 8, 16, 32, 64, and 96 nodes. On each node, AppScale runs a Database Slave/Peer, a Protocol Buffer Server, and an AppServer. Each thread accesses a single AppServer, so nine are in use in our setting. For each configuration, there is also a head node that implements the AppController and the Database Master if there is one (otherwise it is a Database Peer).

The second set of experiments test the performance of the primitive operations of the system when performed as part of an overall workload. Here, the number of nodes is constrained to 16 nodes (due to space limitations) and 10000 random operations are performed with a 50/30/20 get/put/query ratio. Once an operation is selected, nine concurrent threads perform the operation and access their corresponding AppServers.

4.2 Results

We first present results for web application response time between Active Cloud DB and the datastores. Response time includes the round-trip time between Active Cloud DB and datastore including the processing overhead of the Protocol Buffer Server.

Figure 1 displays results for the get, put, and delete operations. The left graph shows the performance of the get operation. The additional entry points for Cassandra allows it to process reads faster than MemcacheDB. Varying the number of nodes in the system does not have a significant impact on the performance of the get operation. With caching, there is an improvement in performance for both Cassandra and MemcacheDB. This is because write-through caching leads to cache hits for all




Figure 1: Average round-trip time for get (left), put (middle), and delete (right) operations under a load of 9 concurrent threads for a range of different AppScale cloud sizes (node count).




Figure 2: Average time for query operation for different node configurations and 9 concurrent threads. The legend is the same as that in Figure 1.

but the first event and cache access is significantly faster than datastore access.

The middle graph shows the performance of the put operation. As was the case in [2], the increased number of entry points for Cassandra allow it to process writes faster than MemcacheDB. Like in the case of the get operation, we see that the performance does not drastically change in either direction for either database with respect to the number of nodes in the system. We see the same trends occurring for the same datastores when caching is employed, and about the same performance for a datastore regardless of whether the cache is employed or not. This shows that the overhead of performing caching is negligible with respect to the overall time of the operation. The right graph shows the performance of the delete operation. Deletes perform similarly to puts for both datastores as well as with and without caching.

Figure 2 shows the performance of the query operation. This operation is to be the slowest in the system since it operates on an entire table and returns all the keys instead of operating on a single key. For our experiments

a query operation returns all 1000 items in the datastore, and we see that having more entry points negatively impacts the datastore’s ability to return all the items for a given table. This impact is consistent across the various node deployments, with Cassandra consistently performing worse than MemcacheDB.

Employing the caching scheme negates this difference since all reads (except for the first) access the cache instead of the datastore. This yields results and conclusions very similar to that for the get operation, but with a degraded performance due to the fact that our caching scheme must marshal and un-marshall the data when caching it. Therefore, un-marshalling all 1000 items in the datastore constitutes the difference between these operations. Similarly to the get operation, this speedup only applies without writes: a single write causes performance to degrade for the next read.

We next consider a second workload. Figure 3 shows the performance of the system under a 50/30/20 get/put/query workload across 16 nodes. All operations perform faster than in the previous experiments, as this workload is performed on an initially empty database. However, the same trends from before are preserved in this workload analysis. Get operations are still faster for Cassandra than MemcacheDB, but now both are significantly slower than their cached equivalents. This is likely due to the substantially smaller amount of data in memcached, allowing for much faster read access. As was the case in the previous experiments, write performance is roughly the same whether or not caching is employed. Finally, query performance is substantially better than in the previous experiments. This is to be expected since the database has substantially less information in it than in the previous experiments. Caching the data has less of an impact here since the non-cached versions perform much better than in the previous experiments.




Figure 3: Average round-trip time for get (left), put (middle), and query (right) operations under a load of 9 concurrent threads for a 50/30/20 get/put/query workload, run over 16 nodes.

Table 1: Average Active Cloud DB time for each operation using Google App Engine (node count unknown) and AppScale (16 nodes).

	Google App Engine Medium Load	AppScale 16 node Cassandra Cache	AppScale 16 node MemcacheDB Cache
put	0.28	0.19	0.40
get	0.24	0.14	0.14
delete	0.28	0.18	0.14
query	2.62	0.26	0.26
no-op	0.18	0.14	0.14

Finally, we experiment with executing Active Cloud DB on Google resources. Table 1 shows the average response time using a load of three threads as opposed to nine for our original first workload (in-order, repeated operation execution). Google continuously killed our 9-thread even though we were within our quota. We also summarize the AppScale results (the same as from the previous graphs using the 16-node configuration). For this data, we include the no-op data which we did not present in graph form. Note that for Google experiments, we have no control over the number of nodes we have been allocated for our application.

5 Conclusion

We present Active Cloud DB, an application-level interface service to the Google App Engine (GAE) cloud. Active Cloud DB is a GAE application that executes over GAE or over its open-source counterpart, AppScale. It exposes the Google datastore API via HTTP to other languages and frameworks. We evaluate its use within Google and AppScale and present a number of proof of concept applications that make use of the interface to access a wide range of diverse key-value stores easily and automatically. We also extend AppScale with simple caching support to significantly improve query performance for Active Cloud DB and other GAE applications that execute using an AppScale cloud. The proof of con-

cept applications, AppScale, and Active Cloud DB, can all be found at <http://appscale.cs.ucsb.edu>.

References

- [1] V. Bala, E. Duesterwald, and S. Banerjia. Dynamo: a transparent dynamic optimization system. *ACM SIGPLAN Notices*, 35(5):1–12, 2000.
- [2] C. Bunch, N. Chohan, C. Krintz, J. Chohan, J. Kupferman, P. Lakhina, Y. Li, and Y. Nomura. Key-Value Datastores Comparison in AppScale. Technical Report 2010-03, UC Santa Barbara, 2010.
- [3] Cassandra. <http://incubator.apache.org/cassandra/>.
- [4] F. Chang, J. Dean, S. Ghemawat, W. Hsieh, D. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. Gruber. Bigtable: A Distributed Storage System for Structured Data. In *Symposium on Operating System Design and Implementation*, 2006.
- [5] N. Chohan, C. Bunch, S. Pang, C. Krintz, N. Mostafa, S. Soman, and R. Wolski. AppScale: Scalable and Open AppEngine Application Development and Deployment. In *International Conference on Cloud Computing*, Oct. 2009.
- [6] B. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears. Benchmarking Cloud Serving Systems with YCSB, Mar. 2010. <http://www.brianfrankcooper.net/pubs/yccb.pdf>.
- [7] MemcacheDB. <http://memcachedb.org/>.
- [8] D. Nurmi, R. Wolski, C. Grzegorzczak, G. Obertelli, S. Soman, L. Youseff, and D. Zagorodnov. The Eucalyptus Open-source Cloud-computing System. In *IEEE International Symposium on Cluster Computing and the Grid*, 2009. <http://open.eucalyptus.com/documents/ccgrid2009.pdf>.
- [9] Protocol Buffers. Google’s Data Interchange Format. <http://code.google.com/p/protobuf/>.
- [10] D. Thomas and D. Hansson. Agile Web Development with Rails: Second Edition. 2006.