

# IBM Research Report

## Universal Data Store Manager

**Arun Iyengar**

IBM Research Division  
Thomas J. Watson Research Center  
P.O. Box 218  
Yorktown Heights, NY 10598 USA



# Universal Data Store Manager<sup>1</sup>

ARUN IYENGAR, IBM Research, aruni@us.ibm.com

---

This paper describes work that we have done in developing infrastructure components for efficient data management. A key component is the use of multiple options for storing data including file systems, SQL (relational) databases, NoSQL databases, and caches. We provide monitoring capabilities so users can determine the performance of different data stores and pick the most appropriate one.

Another key aspect of this work is enhanced storage clients which are particularly applicable to cloud storage. Our work extends existing clients by adding integrated caching support, encryption, compression, and delta encoding. We have built enhanced clients for multiple storage systems including Cloudant and OpenStack Object Storage. We have also written a Java library for allowing our enhanced client features to be integrated with other storage clients and a wide range of applications.

---

## 1. INTRODUCTION

A broad range of data stores are currently available including file systems, SQL (relational) databases, NoSQL databases, caches, etc. An increasing number of data stores are cloud based. A key challenge is providing applications with easy access to a wide range of storage options. We have developed a universal data store manager (UDSM) which gives application programs access to a wide range of storage options. The UDSM also provides monitoring capabilities. That way, users can determine the performance of different storage systems and compare them to pick the best option.

A second aspects of this work is the development of enhanced storage clients. Storage systems typically consist of clients accessing data on one or more storage servers. The clients and servers communicate via a protocol such as HTTP. Widely used storage systems such as Cloudant (built on top of CouchDB), OpenStack Object Storage, and Cassandra have clients which are written for specific programming languages (e.g. Java, Python, JavaScript, etc). These clients handle low level details such as communications with the server using an underlying protocol such as HTTP. That way, client applications can communicate with the storage server via method (or other type of subroutine) calls in the language in which the client is written. Examples of such clients include the Cloudant Java client [Cloudant], the Java library for OpenStack Storage (JOSS) [Javaswift], and the Java Driver for Apache Cassandra[DataStax] (Figure 1).

---

<sup>1</sup> This paper extends Arun Iyengar, *Enhanced Storage Clients*, IBM Research Report RC25584 (WAT1512-042), December 2015,

[http://domino.watson.ibm.com/library/CyberDig.nsf/papers/16214813202B330D85257F2A004A2187/\\$File/rc25584.pdf](http://domino.watson.ibm.com/library/CyberDig.nsf/papers/16214813202B330D85257F2A004A2187/$File/rc25584.pdf)

and Arun Iyengar, *Infrastructure Components for Efficient Data Management*, IBM Research Report RC25599 (WAT1604-013), April 7, 2016,

[http://domino.watson.ibm.com/library/CyberDig.nsf/papers/71C046EA5608285085257F9500647961/\\$File/rc25599.pdf](http://domino.watson.ibm.com/library/CyberDig.nsf/papers/71C046EA5608285085257F9500647961/$File/rc25599.pdf)

## Storage Clients and Servers




Fig. 1. Storage clients are essential for communicating with storage servers.

This paper describes work that we have done in enhancing storage clients to support integrated caching, data compression, encryption, and delta encoding. These features considerably enhance the functionality and performance of storage systems and provide key features that are needed by application programs.

The remainder of this paper is structured as follows. Section 2-3 describe our universal data store manager. Sections 3-7 describe key features of our enhanced storage clients. Section 8 describes how to use open source implementations of the UDSM and enhanced storage clients.

## 2. UNIVERSAL DATA STORE MANAGER (UDSM)

The Universal Data Store Manager (UDSM) allows application programs to access multiple data stores including file systems, SQL (relational) databases, NoSQL databases, and caches (Figure 2). The current implementation has support for file systems, MySQL [MySQL], Cloudant [Cloudant], Object Storage (which implements the OpenStack Swift API) [OpenStack], Redis [Redis], and an in-process cache.

The UDSM provides a common key-value interface. Each storage system implements the key-value interface. That way, it is easy to switch from one storage system to another. The key-value interface exposed to application programs hides the details of how the interface is actually implemented by the underlying storage system.

In some cases, a key-value interface is not sufficient. For example, a MySQL user may need to issue SQL queries to the underlying database. The UDSM allows the user to access native

features of the underlying storage system when needed. That way, applications can use the common key-value interface when appropriate as well as all other capabilities of the underlying storage systems when necessary.

A key feature of the UDSM is the ability to monitor different storage systems for performance. That way, users can measure and compare the performance of different storage systems. The UDSM collects both summary performance statistics such as average latency as well as detailed performance statistics such as historical latency measurements. It is often desirable to only collect latency measurements for recent requests. There is thus the capability to collect detailed data for recent requests while only retaining summary statistics for older data.

Performance data can be stored persistently using any of the storage systems supported by the UDSM.

Both the UDSM and the enhanced storage clients described later in the paper provide data encryption, compression, and serialization. The use of the key-value interface also makes it very easy to provide caching support for storage systems which are not caches. Since all of the storage systems which are caches (e.g. Redis, the in-process cache) implement the same key-value interface, similar method calls can be used to store data in the main storage system as well as in the cache.

## Universal Data Store Manager (UDSM)

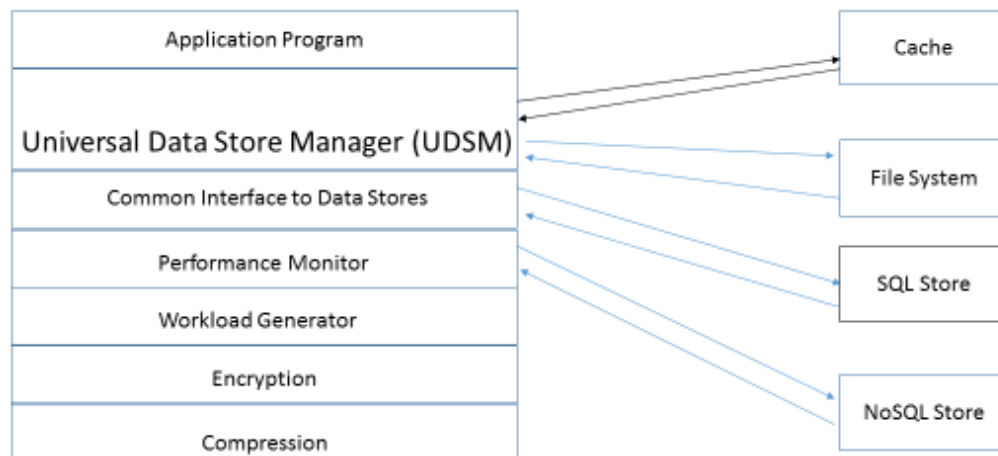


Fig. 2. A Universal Data Store Manager (UDSM).

### 3. PERFORMANCE TESTING

Since storing and retrieving data can have significant overhead, it is important to understand the performance of data storage systems. The UDSM allows storage operations to be monitored so

that users can understand the overhead of their data storage systems as well as compare the overhead of different storage systems. We also provide automatic performance testing features which allow users to quickly and easily test the performance of different data stores under different workloads.

In order to test and compare performance of different data stores, we provide a workload generator which generates storage requests and determines the time for satisfying the requests. The workload generator can synthetically generate data objects to be stored. Alternatively, users can provide their own data objects for performance tests either by placing the data in input files or writing a user-defined method to provide the data.

Our automated performance tester determines performance both with and without caching using either an in-process or remote process cache. Performance is determined for a variety of different hit rates, where the granularity of hit rates is specified by users.

Performance for encryption, decryption, compression, and decompression is also provided.

Data from performance testing is stored in text files which can be easily imported into graph plotting tools such as gnuplot, spreadsheets such as Microsoft Excel, and data analysis tools such as MATLAB.

A key aspect of our performance tester is the ease of generating performance data. Any client using our libraries can easily generate performance data across all data stores supported by our libraries.

#### 4. ENHANCED STORAGE CLIENTS

Our enhanced storage clients are built on top of a storage client library which handles features such as caching, encryption, compression, and delta encoding (Figure 3). For each of these features, there is an interface and multiple possible implementations. For example, there are multiple caching implementations and multiple encryption implementations which a storage client can choose from. The storage client library (SCL) is available as a standalone entity. Storage clients for a specific storage system such as Cloudant, OpenStack Object Storage, Cassandra, etc. can be integrated with the storage client library (SCL) by adding SCL API calls at critical points in the client code. For example, when a storage client queries the server for an object, an SCL API call could be inserted to first look for the object in the cache. When a storage client updates an object at the server, an SCL API call could be inserted to update (or invalidate) the object in the cache.

If these SCL API calls are integrated with the storage client, this results in an enhanced storage client with considerable performance enhancements and added functionality over a vanilla client. An alternative approach which does not require changes to storage clients is to import the client library directly into an application. This allows applications to directly use SCL features such as caching, encryption, and compression in the most appropriate way.

## Enhanced Data Store Client

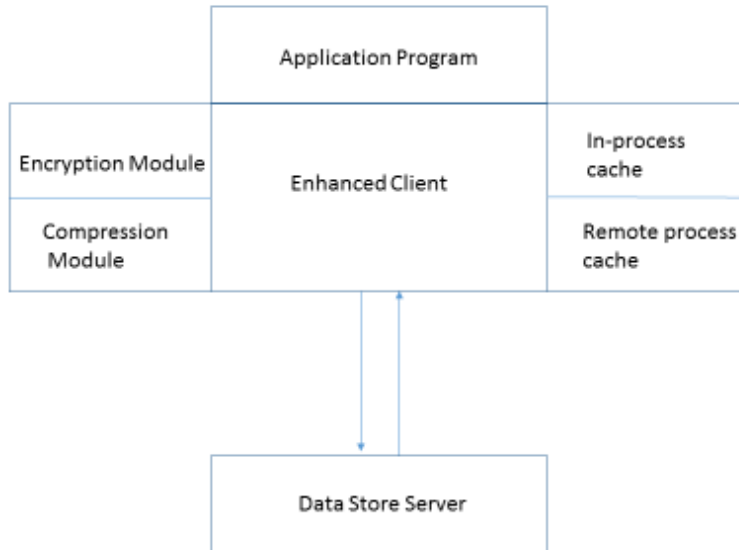


Fig. 3. An enhanced data store client.

### 5. CACHING

Caching support is critically important for improving performance [Iyengar 97]. The latency for communicating between clients and servers can be high. Caching can dramatically reduce this latency. If the cache is properly managed, it can also allow a client program to continue executing in the presence of poor or limited connectivity. Caching is most effective for data which does not change, or which is only updated rarely. Our caches have an API which controls which data are cached and allows caches to be updated. Applications can add and delete data from caches. Furthermore, cached data can have expiration times.

A key feature of our architecture is that it is modular. Our SCL includes a Cache interface which defines how an application interacts with the cache. There are multiple implementations of the Cache interface which applications can choose from (Figure 4).

There are two types of caches. In-process caches store data within the process corresponding to the application (Figure 5) [Iyengar 99]. That way, there is no interprocess communication required for storing the data. For our Java implementations of in-process caches, Java objects can directly be cached. Data serialization is not required. In order to reduce overhead when the object is cached, the object (or a reference to it) can be stored directly in the cache. This means that changes to the object from the application could affect changes to the cached object itself. In order to prevent the value of a cached object from being modified by changes to the object being

made in the application, a copy of the object can be made before the object is cached. This results in overhead for copying the object.

Another approach is to use a remote process cache (Figure 6) [Iyengar 97]. In this approach, the cache runs in one or more separate processes from the application. A remote process cache can run on a separate node from the application as well. There is some overhead for communication with a remote process cache. In addition, data often has to be serialized before being cached. However, remote process caches also have some advantages over in-process caches. A remote process cache can be shared by multiple clients, and this feature is often desirable. Remote process caches can often be scaled across multiple processes and nodes to handle high request rates and increase availability.

## Multiple Implementations of Cache Interface

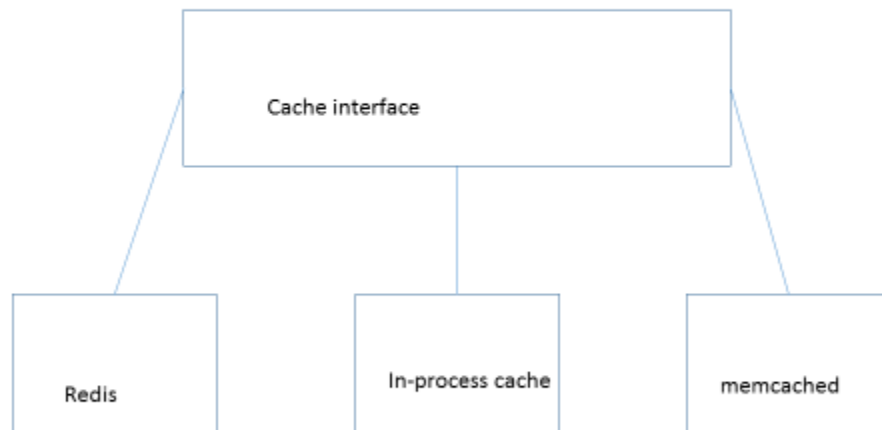



Fig. 4. Our SCL supports multiple cache implementations.


## In-Process Cache



Extremely fast  
Cached objects don't have to be serialized  
Cache not shared by multiple clients

Fig. 5. In-process caches store data within the application process.

## Remote Process Cache



Multiple clients can share cache  
Cache can scale to many processes  
Overhead for interprocess communication  
Cached objects may need to be serialized

Fig. 6. Remote process caches store data outside of the application process.



There are several caches available that are available as open source solutions. Redis [Redis] and memcached [Memcached] are widely used remote process caches. They can be used for storing serialized data across a wide range of languages. Clients for redis and memcached are available across a wide range of languages (e.g. Java, C, C++, Python, Javascript, PHP, several others).

Examples of caches targeted at Java environments include the Guava cache [Decker], Ehcache [Ehcache], and OSCache [OSCache]. A common approach is to use a data structure such as a HashMap or a ConcurrentHashMap with features for thread safety and cache replacement. Since there are several good open source alternatives available, it is probably better to use an existing cache implementation instead of writing another cache implementation unless the user has specialized requirements not handled by existing caches.

Our SCL allows any of these caches to be plugged into its modular architecture. In order to use one of these caches, an implementation of the SCL Cache interface needs to be written for the cache. We have already implemented SCL Cache interfaces for redis and the Guava cache and are working on interfaces for other caches such as memcached.

The SCL allows applications to assign (optional) expiration times to cached objects. After the expiration time for an object has elapsed, the cached object is obsolete and should not be returned to an application until the server has been contacted to either provide an updated version or verify that the expired object is still valid. Cache expiration times are managed by the SCL and not by the underlying cache. There are a couple of reasons for this. Not all caches support expiration times. A cache which does not handle expiration times can still implement the SCL Cache interface. In addition, for caches which support expiration times, objects whose expiration times have elapsed might be purged from the cache. We do not always want this to happen. After the expiration time for a cached object has elapsed, it does not necessarily mean that the object is obsolete. Therefore, the SCL has the ability to keep around a cached object *o1* whose expiration time has elapsed. If *o1* is requested after its expiration time has passed, then the client might have the ability to revalidate *o1* in a manner similar to an HTTP GET request with an If-Modified-Since header. The basic idea is that the client sends a request to fetch *o1* only if its version of *o1* has changed. In order to determine if its version of *o1* is obsolete, the client could send a timestamp, entity tag, or other information identifying the version of *o1* stored at the client. If the server determines that the client has an obsolete version of *o1*, then the server will send a new version of *o1* to the client. If the server determines that the client has a current version of *o1*, then the server will indicate that the version of *o1* is current (Figure 7).

## Handling Expiration Times




Fig. 7. Revalidation with the server after expiration times have expired can avoid fetching cached objects which have not changed, even if their expiration times have passed.

Using this approach, the client does not have to receive identical copies of objects whose expiration times have elapsed even though they are still current. This can save considerable bandwidth and improve performance. There is still latency for revalidating `o1` with the server, however.

If caches become full, a cache replacement algorithm such as least recently used (LRU) or greedy-dual-size [Cao] can be used to determine which objects to retain in the cache.

Caching can be used for situations in which there is poor connectivity between the client and server. In some cases, an application can be written so that it can execute for considerable periods of time using locally cached data with sporadic batch updates with the server.

Some caches such as redis have the ability to back up data in persistent storage (e.g. to a hard disk or solid-state disk). This allows data to be preserved in the event that a cache fails. It is also often desirable to store some data from a cache persistently before shutting down a cache process. That way, when the cache is restarted, it can quickly be brought to a warm state by reading in the data previously stored persistently.

Data can be replicated across multiple caches as another alternative to preserving cached data in the event of a cache failure. This approach can slowly be used to scale caches to handle high request rates. Multiple data copies introduce the problem of data inconsistencies. Caching generally works best for data which does not change very frequently.

## 6. ENCRYPTION AND DATA CONFIDENTIALITY

Data confidentiality is critically important. Many storage systems encrypt data within the server. However, this may not be sufficient. Users of a storage system might prefer to encrypt data before it reaches the server since the server might not be completely trustworthy. It may also be undesirable to store confidential information in a cache as the contents of the cache could be read by a malicious third party. Both the data in a cached object as well as the key used to index the object in a cache could contain confidential information.

Our SCL can encrypt data both before sending data to the server as well as before caching the data. The architecture supports different encryption algorithms. We currently are using an implementation of the Advanced Encryption Standard (AES) [NIST].

There are situations in which data anonymization is needed but not strong encryption. In this case, a data anonymization implementation of the encryption module can be used. The data anonymizer might preserve anonymity while still allowing some types of computations over the anonymized data. Since decryption would not be needed for analyzing the data, performance might be improved. In addition, privacy would be better preserved since the actual data would not be fully exposed, even during data analysis.

## 7. COMPRESSION AND DELTA ENCODING

Compression can considerably reduce the space consumed by data objects. Even if the storage server provides compression, compressing data at the client may be desirable to reduce data transfer sizes between the client and server. Furthermore, it may be desirable to compress objects before caching them in order to conserve cache space. The SCL has the capability to compress data using gzip [Gzip]. Other compression algorithms can also be used.

Data transfer sizes between the client and server can further be reduced by delta encoding. The key idea is that when the client updates an object *o1*, it may not have to send the entire updated copy of *o1* to the server. Instead, it sends a delta between *o1* and the previous version of *o1* stored at the server. This delta might only be a fraction of the size of *o1* [Douglass].

A simple example of delta encoding is shown in Figure 8. Only two elements of the array in the figure change. Instead of sending a copy of the entire updated array, only the delta shown at the bottom is sent. The first element of the delta indicates that the 5 array elements beginning with index 0 are unchanged. The next element of the delta contains updated values for the next two consecutive elements of the array. The last element of the delta indicates that the 6 array elements beginning with index 7 are unchanged.

## Simple delta encoding example




Fig. 8. Example of delta encoding.

Our delta encoding algorithm uses key ideas from the Rabin-Karp string matching algorithm [Karp]. Data objects are serialized to byte arrays. Byte arrays can be compressed by finding substrings previously encountered. If the server has a previous substring, the client can send bytes corresponding to the substring by sending an index corresponding to the position of the substring and the length of the substring as illustrated in Figure 6. Matching substrings should have a minimum length, `WINDOW_SIZE` (e.g. 5). If the algorithm tries to encode differences by locating very short substrings (e.g. of length 2), the space overhead for encoding the delta may be higher than simply sending the bytes unencoded. When a matching substring of length at least `WINDOW_SIZE` is found, it is expanded to the maximum possible size before being encoded as a delta.

As we mentioned, a cached object `o` can be serialized to a byte array, `b`. We find matching substrings of `o` by hashing all subarrays of `b` of length `WINDOW_SIZE` in a hash table. In order to hash substrings of `o` efficiently, we use a rolling hash function which can be efficiently computed using a sliding window moving along `b`. That way, the hash value for the substring beginning at `b[i + 1]` is efficiently calculated from the hash value for the substring beginning at `b[i]`.

## 8. IMPLEMENTATION

### 8.1 Using the Uniform Data Store Manager (UDSM)

We have implemented a UDSM which allows Java applications to access multiple data stores and provides caching, encryption, and compression. The UDSM is available as open source software on github [Storage-manager]. This UDSM can be used for enhancing data store clients as well as for a wide variety of other applications when access to multiple data stores, performance monitoring and comparison of data stores, caching, encryption, and/or compression are needed.

We now give some examples of how the UDSM can be used. A common key-value interface is defined in `com.ibm.storage.storagemanager.interfaces.KeyValue<K, V>`. This interface is implemented by multiple data stores. That way, an application program which uses this interface can easily switch between different data store implementations. In order to use a file system implementation of the key-value interface, the following can be used:

```
import com.ibm.storage.storagemanager.implementations.file.KeyValueFile;
import com.ibm.storage.storagemanager.interfaces.KeyValue;
```

```
    KeyValue<String, Integer> datastore = new KeyValueFile<String, Integer>("db1", true);
```

In this example, the keys are strings, and the values are integers. Each value will be stored in a file whose name is the key. The files will be in the directory `database/db1`. The last parameter, “true”, indicates that all previous files stored in this directory should be deleted.

In order to use an implementation of the key-value interface in which a cache is used which runs in the same process as the application, the following can be used:

```
import com.ibm.storage.storagemanager.implementations.guava.KeyValueGuava;
import com.ibm.storage.storagemanager.interfaces.KeyValue;
```

```
    KeyValue<String, Integer> datastore = new KeyValueGuava<String, Integer>(numObjects);
```

where `numObjects` is the maximum number of objects which can be stored in the cache.

In order to use an implementation of the key-value interface in which a Redis [Redis] cache is used, a Redis cache needs to be running in a separate process. The following can then be used:

```
import com.ibm.storage.storagemanager.implementations.redis.KeyValueRedis;
import com.ibm.storage.storagemanager.interfaces.KeyValue;
```

```
    KeyValue<String, Integer> datastore = new KeyValueRedis<String, Integer>("localhost",
6379, 60);
```

In this example, the cache is running on the same node as the application program, 6379 is the port number, and 60 indicates that idle connections should be closed after 60 seconds.

In order to use an implementation of the key-value interface in which a MySQL [MySQL] database is used,

```
import com.ibm.storage.storagemanager.implementations.sql.KeyValueSQL;
import com.ibm.storage.storagemanager.interfaces.KeyValue;
```

```
    KeyValue<String, Integer> datastore = new KeyValueSQL<String, Integer>("db1",
CONFIG_FILE, true);
```

In this example, “db1” is the name of the database for storing the key-value data, `CONFIG_FILE` is a configuration text file, and `true` indicates that all previous values stored in the database should be deleted. A skeleton for the configuration file is in the gitub repository file `config/sql.config`. The first line is the name of the sql database for storing the data. The second line is the url for accessing the database. The third line is the user id, and the fourth line is the password.

In order to pass parameters directly to the `KeyValueSQL` constructor and avoid using a configuration file, the following constructor can be invoked:

```
/**
 * Constructor. Establishes a session with a Key-value store..
 *
 * @param tableName
 *         identifies the table name for the key-value pairs
 * @param databaseName
 *         database name
 * @param url
 *         base URL for the data service
 * @param userName
 *         user name for authentication
 * @param passWord
 *         password
 */
public KeyValueSQL(String tblName, String databaseName, String url, String userName,
String passWord,
    boolean clearAll);
```

In order to use an implementation of the `KeyValue` interface in which `Cloudant` [`Cloudant`] is used,

```
import com.ibm.storage.storagemanager.implementations.cloudant.KeyValueCloudant;
import com.ibm.storage.storagemanager.interfaces.KeyValue;
KeyValue<String, Integer> datastore = new KeyValueCloudant<String, Integer>("db1",
CONFIG_FILE, true);
```

In this example, “db1” is the name of the database for storing the key-value data, `CONFIG_FILE` is a configuration text file, and `true` indicates that all previous values stored in the database should be deleted. A skeleton for the configuration file is in the gitub repository file `config/cloudant.config`. The first line is the url for the `Cloudant` account. The second line is the user id, and the third line is the password.

In order to pass parameters directly to the `KeyValueCloudant` constructor and avoid using a configuration file, the following constructor can be invoked:

```
/**
 * Constructor. Establishes a session with a Key-value store.
 *
 * @param storeId
 *         identifies the store, could be a URL such as for cloudant
```

```

* @param dbName
*     identifies the database
* @param userId
*     identifies the user
* @param password
*     password corresponding to the user
* @param deletePreviousDb
*     true if previous version of dbName should be deleted
*/
public KeyValueCloudant(String storeId, String dbName, String userId, String password,
    boolean deletePreviousDb)

```

In order to use an implementation of the `KeyValue` interface in which OpenStack Object Storage [OpenStack] is used,

```

import com.ibm.storage.storagemanager.implementations.objectstorage.KeyValueObjectStorage;
import com.ibm.storage.storagemanager.interfaces.KeyValue;
    KeyValue<String, Integer> dataStore = new KeyValueObjectStorage<String, Integer>("db1",
    CONFIG_FILE, true);

```

In this example, “db1” is the name of the database (container) for storing the key-value data, `CONFIG_FILE` is a configuration text file, and `true` indicates that all previous values stored in the database should be deleted. A skeleton for the configuration file is in the gitub repository file `config/objectstorage.config`. The first line is the url for accessing Object Storage. The second line is the user id, and the third line is the password.

In order to pass parameters directly to the `KeyValueObjectStorage` constructor and avoid using a configuration file, the following constructor can be invoked:

```

/**
 * Constructor. Establishes a session with a Key-value store.
 *
 * @param dbName
 *     identifies the database
 * @param storeId
 *     Url for the data store
 * @param userId
 *     user ID
 * @param password
 *     password
 * @param clearAll
 *     true if all previous database entries should be deleted
 */
public KeyValueObjectStorage(String dbName, String storeId, String userId, String password,
    boolean clearAll)

```

The same `KeyValue` methods are used for each implementation of `datastore`. This makes it easy to substitute different implementations of data stores within an application program. We next show examples of using `datastore`. Note that the same method calls could be used with `KeyValueCloudant`, `KeyValueFile`, `KeyValueGuava`, `KeyValueRedis`, `KeyValueSQL`, `KeyValueObjectStorage`, and any other classes which implement the `KeyValue` interface.

The following method call adds 42 to the data store indexed by `key1`:

```
datastore.put(key1, 42);
```

The following method call adds all key-value pairs corresponding to `map` to the data store and returns the number of key-value pairs successfully stored:

```
int numStored = datastore.putAll(map);
```

The following method call returns the value corresponding to `key`. It returns null if `key` is not found in the data store.

```
val = datastore.get(key);
```

In the following method call, `list` is a list of keys. `getAll` looks up all key-value pairs corresponding to keys in `list` and returns a map containing them. In this example, keys are strings, and values are integers.

```
Map<String, Integer> map = datastore.getAll(list);
```

The following method call deletes the key-value pair corresponding to `key2` from the data store if present:

```
datastore.delete(key2);
```

In the following method call, `list` is a list of keys. `deleteAll` deletes all key-value pairs corresponding to a key in `list`. The number of deleted key-value pairs is returned:

```
numDeleted = datastore.deleteAll(list);
```

The following method call deletes all objects in the data store:

```
datastore.clear();
```

The following method call outputs a string identifying the type of data store represented by `datastore`:

```
datastore.storeType();
```

The `size` method returns the number of objects in the data store:

```
System.out.println("Data store size: " + datastore.size());
```

The following displays the contents of the entire data store. It should not be invoked if the data store contains a large amount of data as the data outputted would be prohibitively large:

```
System.out.println(datastore.toString());
```



In order to monitor the performance of a data store, the following import statements should be used:

```
import com.ibm.storage.storagemanager.implementations.monitor.MonitoredKeyValue;
import com.ibm.storage.storagemanager.interfaces.KeyValue;
```

The following creates an object which is used to monitor data stores; *datastore* is of type *KeyValue* (which includes any implementation of the *KeyValue* interface):

```
MonitoredKeyValue<String, Integer> monitoredStore = new MonitoredKeyValue<String,
Integer>(datastore, 10);
```

where 10 is the number of most recent data points to keep around for each transaction type. The *MonitoredKeyValue* class implements an interface which extends the *KeyValue* interface. Thus, the same methods used for the *KeyValue* interface, such as *put* and *get* described earlier, can also be applied to *monitoredStore*. The key point is that when this is done, *monitoredStore* records performance statistics on the latency for executing those methods. The methods for accessing, setting, and clearing all of the performance stats are defined in:

<https://github.com/aruniyengar/storage-manager/blob/master/src/main/java/com/ibm/storage/storagemanager/interfaces/KeyValueMonitor.java>

After *monitoredStore* has been used for accessing a data store, the performance statistics from the data store can be obtained via:

```
import com.ibm.storage.storagemanager.implementations.monitor.RequestStats;
import com.ibm.storage.storagemanager.implementations.monitor.StorageStats;
```

```
StorageStats stats = monitoredStore.getStorageStats();
System.out.println(stats.allStats());
System.out.println("Total number of requests: " + stats.getNumRequests());
System.out.println("Total time spent on storage: " + stats.getTotalRequestTime());
System.out.println("Start time: " + stats.getStartTime());
System.out.println("End time: " + stats.getEndTime());
stats.setEndTimeNow();
System.out.println("New End time: " + stats.getEndTime());
System.out.println("Store type: " + stats.getStoreType());
RequestStats stats2 = stats.getRequestData(RequestType.PUT);
System.out.println("Statistics for request type: " + stats2.getRequestType());
System.out.println("Number of requests: " + stats2.getNumRequests());
System.out.println("Time taken by requests: " + stats2.getTotalRequestTime());
System.out.println("Recent request times: " +
Arrays.toString(stats2.getRecentRequestTimes()));
```

The UDSM also includes a workload generator which is used to measure and compare the performance of different data stores. The performance of different data stores can be compared using the specified workload. The following can be used to determine the performance of *datastore*:

```
import com.ibm.storage.storagemanager.performancetester.PerformanceTester;
PerformanceTester pt = new PerformanceTester(configFile);
```

`pt.testDataStore(datastore);` // `datastore` is of type `KeyValue<String, byte[]>`  
`configFile` is a string indicating the configuration file specifying the directory for the output files and the parameters for the workload generator. A skeleton for the configuration file is contained in:

<https://github.com/aruniyengar/storage-manager/blob/master/config/performanceconfig>

The workload generator outputs performance data in text files which can be visualized using tools such as gnuplot, Excel, and MATLAB.

In order to run the workload generator on multiple data stores using a single method call, the following static method from `com.ibm.storage.storagemanager.performancetester.PerformanceTester` can be used:

```
/**
 * Run all tests, store the results in output files.
 *
 * @param configFile
 *       file storing configuration parameters
 * @param dataStores
 *       Array containing all data stores to test
 */
public static void runAllTests(String configFile, ArrayList<KeyValue<String, byte[]>>
dataStores);
```

There is a JUnit test to easily run performance tests across all data stores currently supported. This JUnit test is performed by the method

```
public void allTests();
```

in the Java class:

<https://github.com/aruniyengar/storage-manager/blob/master/src/test/java/com/ibm/storage/storagemanager/performancetester/tests/PerformanceTests.java>

When Redis is being used, it may be desirable to use features of Redis which go beyond the methods offered by the `KeyValue` interface. The following `KeyValueRedis` method returns a data structure corresponding to the Jedis interface for Redis [`Jedis`] which allows application programs to access the cache using Jedis methods:

```
/**
 * Return underlying Jedis object for applications to explicitly use.
 *
 * @return value underlying Jedis object representing cache
 *
 */
public Jedis getDatabase();
```

When an SQL store is being used, it may be desirable to issue SQL queries which go beyond the methods offered by the `KeyValue` interface. The following `KeyValueSQL` method

returns a connection data structure which allows application programs to make SQL queries using JDBC:

```
/**
 * Return database connection for applications to explicitly use.
 *
 * @return connection to database
 *
 * */
public Connection getConnection();
```

When Cloudant is being used, it may be desirable to access Cloudant in ways which go beyond the methods offered by the Key Value interface. The following Key Value Cloudant method returns a data structure which allows application programs to access Cloudant using the Cloudant Java client:

```
/**
 * Return underlying Database object for applications to explicitly use.
 *
 * @return value underlying Database object representing database
 *
 * */
public Database getDatabase();
```

When Object Storage is being used, it may be desirable to access Object Storage in ways which go beyond the methods offered by the Key Value interface. The Java library for OpenStack Storage (JOSS) [Jawaswift] provides a convenient Java interface for using Object Storage. The following Key Value Object Storage method returns a Container object which allows application programs to access the data store using JOSS:

```
/**
 * Return underlying container for applications to explicitly use.
 *
 * @return value underlying Container object representing database
 *
 * */
public Container getDatabase();
```

In order to use encryption and decryption, the following class should be imported:  
**import** com.ibm.storage.storagemanager.util.Encryption;

The following method call generates an encryption key for encrypting data:  
 Encryption.Key secretKey = Encryption.generateKey();

The following method call encrypts *hm* using encryption key *secretKey*:  
 SealedObject so = Encryption.encrypt(hm, secretKey);

The following method call decrypts *so* using encryption key *secretKey*:

```
HashMap<String, Integer> hm2 = Encryption.decrypt(so, secretKey);
```

In order to use compression and decompression, the following class should be imported:

```
import com.ibm.storage.storagemanager.util.Compression;
```

The following method call compresses “hm”:

```
byte[] compressed = Util.compress(hm);
```

The following method call decompresses “compressed”:

```
HashMap<String, Integer> hm2 = Util.decompress(compressed);
```

## 8.2 Using Enhanced Storage Clients

We have implemented enhanced storage clients for multiple storage systems including Cloudant and OpenStack Object Storage. A library (SCL) implementing several of the features discussed previously is available as open source software on github [Storage-client-library]. The github repository supports caching, encryption, and data compression. We have not yet added delta encoding to the public repository but may do so in the future. This library (SCL) can be used for enhancing storage clients as well as for a wide variety of other applications when caching, encryption, and/or compression are needed.

We now give some examples of how our SCL can be used. In order to use an in-process cache with lifetimes, the following import statements should be used:

```
import com.ibm.storage.clientlibrary.CacheWithLifetimes;
```

```
import com.ibm.storage.clientlibrary.InProcessCache;
```

One can create an in-process cache of integers indexed by strings using the Guava cache with a maximum cache size of “numObjects” and a default lifetime of “defaultLifetime” in milliseconds via:

```
CacheWithLifetimes<String, Integer> cache = new InProcessCache<String, Integer>(
    numObjects, defaultLifetime);
```

In order to use a remote process cache, Redis is a preferred implementation. The following import statements should be used:

```
import com.ibm.storage.clientlibrary.CacheWithLifetimes;
```

```
import com.ibm.storage.clientlibrary.RedisCacheExtended;
```

One can connect to a redis cache of byte arrays indexed by strings running on the same node from a Java program via:

```
CacheWithLifetimes<String, byte[]> cache = new RedisCacheExtended<String, byte[]>
    ("localhost", 6379, 60, defaultExpiration);
```

Here, 6379 is the port number, and 60 indicates that idle connections should be closed after 60 seconds.

The following method call adds 42 to the cache indexed by “key1”. “lifetime” is the lifetime of the cached value in milliseconds:

```
cache.put(key1, 42, lifetime);
```

The following method call adds all key-value pairs corresponding to “map” to the cache and assigns each key-value pair a lifetime of “lifetime” milliseconds:

```
cache.putAll(map, lifetime);
```

The following method call returns the value corresponding to “key2”. It returns null if “key2” is not found in the cache, or if the value is expired.

```
val = cache.get(key2);
```

In the following method call, “list” is a list of keys. `getAll` looks up all key-value pairs corresponding to keys in “list” and returns a map containing them. In this example, cache keys are strings, and values are integers.

```
Map<String, Integer> map = cache.getAll(list);
```

The following method call deletes the key-value pair corresponding to “key2” from the cache if present:

```
cache.delete(key2);
```

In the following method call, “list” is a list of keys. `deleteAll` deletes all key-value pairs corresponding to a key in “list”:

```
cache.deleteAll(list);
```

The following method call deletes all objects in the cache:

```
cache.clear();
```

The size method returns the number of cached objects:

```
System.out.println("Cache size: " + cache.size());
```

The following displays the contents of the entire cache. It should not be invoked if the cache contains a large amount of data as the data outputted would be prohibitively large:

```
System.out.println(cache.toString());
```

The following displays cache statistics, such as hit rates:

```
System.out.println(cache.getStatistics().getStats());
```

When Redis is being used as a remote process cache, it may be desirable to use features of Redis which go beyond the methods offered by the `CacheWithLifetimes` interface. The following `RedisCacheExtended` method returns a data structure corresponding to the Jedis interface for Redis [`Jedis`] which allows application programs to access the cache using Jedis methods:

```
/**
```

```
 * Return underlying Jedis object for applications to explicitly use.
```

```
*/
```

```

* @return value underlying Jedis object representing cache
*
* */
public Jedis getJedis();

```

In order to use encryption and decryption, the following class should be imported:  
**import** com.ibm.storage.clientlibrary.Encryption;

The following method call generates an encryption key for encrypting data:  
 Encryption.Key secretKey = Encryption.generateKey();

The following method call encrypts “hm” using encryption key “secretKey”:  
 SealedObject so = Encryption.encrypt(hm, secretKey);

The following method call decrypts “so” using encryption key “secretKey”:  
 HashMap<String, Integer> hm2 = Encryption.decrypt(so, secretKey);

In order to use compression and decompression, the following class should be imported:  
**import** com.ibm.storage.clientlibrary.Util;

The following method call compresses “hm”:  
 byte[] compressed = Util.compress(hm);

The following method call decompresses “compressed”:  
 HashMap<String, Integer> hm2 = Util.decompress(compressed);

Enhanced storage clients are written by adding SCL API calls at critical points in the source code. That way, users of the enhanced storage clients can get the benefits of caching without having to explicitly add and delete objects from the cache. Our enhanced storage clients handle caching transparently to users.

In some cases, users may want explicit control over the cache contents. This may be important in situations in which cached objects are changing and there is no way for the storage client to make the correct caching decisions to achieve both optimal performance and cache consistency. In these situations, the user can directly use SCL API calls to explicitly control the contents of caches.

The decision to compress or encrypt data often needs to be made at the application level. Therefore, users will often directly use SCL API calls to encrypt or compress specific objects without relying on storage clients to do so. It is also possible for storage clients to define encryption and compression policies wherein certain objects by default will be encrypted and/or compressed before they are sent to the client. Similar policies can be defined to automatically encrypt and/or compress certain objects by default before caching them.

The performance improvements we have achieved with caching can be significant. For data stored remotely in the cloud, we have seen latency improvements of multiple orders of magnitude when caching is used.

## ACKNOWLEDGMENTS

The author would like to thank Rich Ellis and Mike Rhodes for their assistance and support in developing caching support for Cloudant.

## REFERENCES

- Dulcardo Arteaga, Ming Zhao, "Client-side Flash Caching for Cloud Systems", Client-side Flash Caching for Cloud Systems", Proceedings of ACM SYSTOR 2014.
- Pei Cao, Sandy Irani, "Cost-Aware WWW Proxy Caching Algorithms", Proceedings of USITS '97.
- Feng Chen, Michael Mesnier, Scott Hahn, "Client-aware Cloud Storage". Proceedings of the 30th International Conference on Massive Storage Systems and Technology (MSST 2014).
- Cloudant, Cloudant Java Client, <https://github.com/cloudant/java-cloudant>
- DataStax. Java Driver for Apache Cassandra, <http://docs.datastax.com/en/developer/java-driver/2.0/java-driver/whatsNew2.html>
- Colin Decker, "CachesExplained", <https://github.com/google/guava/wiki/CachesExplained>
- Fred Douglass, Arun Iyengar, "Application-specific Delta-encoding via Resemblance Detection", Proceedings of the USENIX 2003 Annual Technical Conference.
- Ehcache, <http://www.ehcache.org/>
- Gzip, "The gzip home page", <http://www.gzip.org/>
- David A. Holland, Elaine Angelino, Gideon Wald, Margo I. Seltzer, "Flash Caching on the Storage Client", Proceedings of the 2013 USENIX Annual Technical Conference.
- Arun Iyengar and Jim Challenger, "Improving Web Server Performance by Caching Dynamic Data", In Proceedings of the USENIX 1997 Symposium on Internet Technologies and Systems (USITS '97), Monterey, CA, December 1997.
- Arun Iyengar, "Design and Performance of a General-Purpose Software Cache". In Proceedings of the 18th IEEE International Performance, Computing, and Communications Conference (IPCCC'99), Phoenix/Scottsdale, Arizona, February 1999.
- Javaswift, JOSS: Java library for OpenStack Storage, aka Swift, <http://joss.javaswift.org/>
- Jedis, "A blazingly small and sane redis java client", <https://github.com/xetorthio/jedis>
- Richard M. Karp, Michael O. Rabin, "Efficient randomized pattern-matching algorithms", IBM Journal of Research and Development, Vol. 31 no. 2, March 1987, pp. 249-260.
- Memcached, <http://memcached.org/>
- MySQL, <https://www.mysql.com/>
- NIST, "Announcing the Advanced Encryption Standard (AES)", Federal Information Processing Standards Publication 197, November 26, 2001, <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>
- OpenStack, "Welcome to Swift's documentation!", <http://docs.openstack.org/developer/swift/>
- OSCache, <https://java.net/projects/oscache>
- Redis, <http://redis.io/>
- Storage-client-library, <https://github.com/aruniyengar/storage-client-library>
- Storage-manager, <https://github.com/aruniyengar/storage-manager>