

The Oracle Cloud Storage Archive for Long-term Storage and Preservation

Pyounguk Cho
Oracle
200 Oracle Parkway
Redwood City, CA 94065
650-506-5297
pyounguk.cho@oracle.com

Art Pasquinelli
Oracle
500 Oracle Parkway
Redwood City, CA 94065
650-607-0035
art.pasquinelli@oracle.com

ABSTRACT

Oracle Public Cloud offers data backup and archive services with a high level of data availability and at an affordable cost. This demonstration highlights key capabilities for direct usage of Oracle Storage and Archive Services via Openstack Swift APIs and NFS interfaces. In addition, it also showcases integrated solutions using storage management tools such as iRODS and CommVault. A review of how both the technical and economic features of Archival Cloud Computing can be employed in new preservation infrastructure modeling, use cases, and academic and business scenarios will be given. The Oracle Archival Cloud offers new infrastructure opportunities to institutions.

1. INTRODUCTION

Oracle Storage Cloud Service is a secure, elastic, reliable, and cost-effective public cloud storage solution. It can be accessed from anywhere, 24/7, and from any device connected to the Internet. With zero investment in hardware, organizations can buy just as much enterprise-grade storage capacity as needed today, and buy more as required.

Oracle Storage Cloud Service provides an easy-to-use solution to store, manage, and consume large amounts of unstructured data over the Internet. Applications can access Oracle Storage Cloud Service programmatically by using either an OpenStack Swift-compatible REST API or Java API. Academic, library, and IT administrators can monitor key storage metrics and manage users and roles by using a web-based graphical console. Users can apply role-based access control for data stored on Oracle Storage Cloud Service at a very granular level. As required, data can be made accessible publicly.

Data that is stored using Oracle Storage Cloud Service is replicated on multiple storage nodes, guaranteeing protection against hardware failure and data corruption. Data is never moved out of the data center without owner permission. Oracle Storage Cloud Service can be employed as a cost-effective, remote backup solution for departmental, library, or enterprise data and applications. By backing up data and applications to Oracle Storage Cloud Service, users can avoid large capital and operating expenditures in acquiring and maintaining

iPres 2015 conference proceedings will be made available under a Creative Commons license.

With the exception of any logos, emblems, trademarks or other nominated third-party images/text, this work is available for re-use under a Creative Commons Attribution 3.0 unported license. Authorship of this work must be attributed. View a [copy of this licence](#).

storage hardware. By automating backup routine to run at scheduled intervals, users can further reduce the operating cost of running a backup process. In the event of a disaster at your site, the data is safe in a remote location, and you can restore it quickly to your production systems. To learn more about Oracle Storage Service and to request for a trial subscription:

cloud.oracle.com/storage.

2. ARCHIVING AND PRESERVATION FEATURES

2.1 Security

Oracle uses enterprise-grade processes and operations to secure your data. For enhanced security, you can use the client-side encryption feature of the Java library. A cycle of encryption and decryption ensures that your data remains secure in the cloud. When customers use the client-side encryption feature of the Java library, for every object that is created in Oracle Storage Cloud Service, a unique symmetric key is generated. The Java library uses this key to encrypt data before storing it. After encrypting client data, the Java library encrypts the symmetric key as an envelope key by using an asymmetric key pair that the client provides. The envelope key is then stored as metadata alongside the object data. When customers use the Java library to access such encrypted objects, the envelope key is first retrieved and decrypted by using the asymmetric key pair that a user provides. The resulting symmetric key is then used to decrypt the object data.

2.2 Data Integrity

When an object is created in Oracle Storage Cloud Service from an uploaded file, the service returns the MD5 checksum of the object. This is in the ETag header of the HTTP response. The client that initiated the backup can verify whether the file was uploaded correctly by comparing the MD5 checksum provided by the service with a locally calculated checksum. Every request to Oracle Storage Cloud Service receives an HTTP response containing a status code, which indicates whether the requested operation was completed successfully. The client that initiated the backup can determine whether the data was backed up reliably, by interpreting the status code returned by Oracle Storage Cloud Service.

2.3 Authentication

Oracle Storage Cloud Service authenticates all requests through an authentication token mechanism. Every request to the service must include a valid authentication token, which the service

provided previously in response to an authentication request containing a valid user name and password. The authentication token expires after 30 minutes.

2.4 Back Up Architecture

To facilitate the efficient and reliable upload of files that are larger than 5 GB, Oracle Storage Cloud Service supports uploading files in segments. This feature is called dynamic large objects. Users can segment a large file into multiple small files, each called a segment and each smaller than 5 GB, and then upload the segments individually to Oracle Storage Cloud Service. Customers must also create a manifest object, which will be used when the objects are downloaded, to concatenate the retrieved segments in the correct sequence and stream them in a single response. Note that customers can use their own convention-based schemes for segmenting large files.

Each operation on Oracle Storage Cloud Service is atomic. It either succeeds completely or fails completely. If the upload of a particular file fails, due to a network problem for example, the file must be uploaded again. Data that was uploaded until the network failure occurred is not saved in the cloud. So before you upload large files, even those that are smaller than 5 GB, consider segmenting them and then uploading the segments individually. With this approach, if the upload of a segment fails, only that segment needs to be uploaded again.

To optimize the storage space used in Oracle Storage Cloud Service, consider compressing data before uploading it. When this is done, data will consume less space in Oracle Storage Cloud Service and will take less time to upload and retrieve. A customer can store multiple directories and files in Oracle Storage Cloud Service with a single request, by packaging and compressing them and uploading the resulting tar.gz or tar.bz2 file.

In Oracle Storage Cloud Service, a container is created for each top-level directory and an object is created for each file.

2.5 New Economic Efficiencies

Oracle Archive Storage Cloud Service provides storage for applications and workloads that require long-term retention at the lowest price in the industry. As a “deep cloud” archive, the Archive Storage Cloud is suited for infrequently accessed large-scale data sets. It is priced at \$0.001 GB/Month which equates to \$12,000/PB/Year.

2.6 References and Citations

- [1] Oracle eBook: Backing Up Data and Applications Securely, Reliably, and Efficiently (September 22, 2015)
Oracle Webpage: <https://cloud.oracle.com/storage>