

An Oracle White Paper
May 2010

Cloud-Stored Offsite Database Backups

Introduction

Cloud Computing allows users to tap into a virtually unlimited pool of computing and storage resources over the Internet (the Cloud). Unlike traditional IT, Cloud users typically have little insight or control over the underlying infrastructure, and they must interact with the computing and storage resources via an Application Programming Interface (API) provided by the Cloud vendors. In exchange for those constraints, Cloud users benefit from utility-like costs, scalability, and reliability, as well as the ability to self-provision resources dynamically and pay only for what they use

The ability to back up Oracle Database in the Cloud is a key part of Oracle's Cloud offering. It allows customers to use Storage Clouds, such as Amazon's Simple Storage Service (S3), as their next-generation offsite backup storage destination. Compared to traditional tape-based offsite storage, Cloud backups are more accessible, faster to restore under most circumstances, and more reliable. Cloud backups are also the right protection for databases running within the compute Cloud.

Why Backup Storage in the Cloud

Good Disaster Recovery (DR) practice requires keeping usable business-critical backups offsite. Organizations have traditionally implemented this by writing backups to tape and shipping the tapes to be stored offsite. This is costly and operationally complex, requiring hardware, personnel, and sound procedures to ensure that the offsite backups are up-to-date, secure, and able to be recalled and used in the face of disaster. While shipping and secure storage are often outsourced, the IT organization of the enterprise retains the burden of ensuring the integrity of the backups and procedures.

The pricing and operational characteristics of Cloud Storage make it a very compelling alternative to shipping tapes offsite. Cloud storage offers pay-as-you-go, elastic self-provisioning, with low prices¹ per unit storage per unit time, making costs easy to predict, control, and map to the workloads of an organization's IT assets. Good Cloud infrastructure offers storage redundancy, security, availability and scalability with geographic distribution that enables it to absorb a broad range of adverse events with minimal or no loss of availability. These characteristics make it an excellent alternative to writing, shipping and storing tapes in a secure location. Last, but not least, backups are created and updated over the network, with minimal or no operator involvement – drastically simplifying operational procedures.

Amazon Web Services (AWS) is the first Cloud vendor that Oracle has partnered with to enable database backup in the Cloud. Simple Storage Service ([S3](#)) is the main storage offering of AWS. S3's simple web-services interface enables applications to store and retrieve any amount of data from anywhere on the Internet. S3 is a highly scalable, reliable, fast, inexpensive data storage infrastructure, and thousands of enterprises small and large rely on it for their production storage needs, from “cold” inexpensive storage to serving rich multimedia in real time to customers worldwide.

Sending backups over the Internet to be stored in the Cloud benefits from the elasticity in capacity and operational expenses typical of Cloud services. It can also simplify your own infrastructure as you no longer need to provide and manage storage (e.g., tapes that need to be rotated, shipped away, etc.).

An important objection to over-the-network Cloud backup is that limited network bandwidths in the public Internet preclude the fast transfer of large data amounts like those of a full backup of a large production database. Amazon Web Services addresses this problem by offering S3 data import and export services, which enable bulk movement of data into and out of S3 by shipping portable disks. For example, after a disaster, S3 can express mail a portable hard drive containing all the backup data for a given database. This makes the cloud storage offering comparable to storing tapes offsite, especially when it is done as part of a complete backup strategy that includes keeping backups onsite as well as offsite.

¹ Storage Cloud benefits from the falling prices of the commodity disks underlying the infrastructure, and the economies of scale of the Cloud operator.

Oracle Secure Backup Cloud Module

The Oracle Secure Backup (OSB) Cloud Module enables an Oracle Database to send its backups to Amazon S3. It is compatible with Oracle Database versions 9i Release 2 and above, and it requires a network connection to the Internet, and provisioning the means of payment to Amazon Web Services.² The Oracle Secure Backup Cloud module can also be used when the database is running within the Amazon Elastic Compute Cloud (EC2), in which case it benefits from the higher internal network bandwidth and no transfer costs into and out of S3.

The Oracle Secure Backup Cloud module is implemented using the Oracle Recovery Manager (RMAN) SBT interface. The SBT interface allows external backup libraries to be seamlessly integrated with RMAN. Consequently, database administrators can continue to use their existing backup tools – Enterprise Manager, RMAN and other scripts, etc. – to perform Cloud backups.

OSB Cloud module is available for Linux 32 and 64, SPARC 64, and Windows 32.

Complete Data Security with Built-in Encryption

Oracle Secure Backup leverages RMAN's ability to encrypt backups to ensure data security. Data security and privacy is particularly important in shared, publicly accessible environments such as the Storage Cloud. While most Storage Cloud vendors provide robust security to ensure that only authorized users can access data, Oracle's encryption of backup data before it leaves your database further mitigates risk of theft or unauthorized access because the backup data remains encrypted both in-transit and at rest in the Cloud.

Compressed Backups for Better Performance

Integration with the Oracle Database engine enables Oracle Secure Backup to identify and skip unused space (blocks) within the database. Users also benefit from RMAN's rich compression capabilities. When transmitting backups over slower networks, such as the public Internet, any reduction in backup size is directly realized as an increase in backup performance.

Database Version Support

The Oracle Secure Backup Cloud Module may be used to back up the following supported versions of Oracle Database: Oracle Database 9i Release 2 or higher, including Oracle Database 11g.

² The Cloud Backup Module is a part of the [Oracle Secure Backup product family](#), and licensed on a per-RMAN channel basis. Oracle Secure Backup is Oracle's next-generation tape backup management solution and it now provides customers the flexibility to back up data to either tape or the Cloud.



Figure 1. Oracle Database backup in the Cloud

Benefits of Oracle Cloud Backup

Oracle's Cloud backup functionality provides advantages over traditional tape-based offsite backups:

- Continuous Accessibility:** Backups stored in the Cloud are always accessible – much in the same way local disk backups are. As such, there is no need to call anyone and no need to ship or load tapes before a restore can be performed. Administrators can initiate restore operations using their standard tools (Enterprise Manager, scripts, etc.) just as if the offsite backup was stored locally. This can help make restores faster and reduce down time from days to hours/minutes in many cases. For large databases where shipping a portable disk from the Cloud is required, a restore takes no longer than it would take to have a tape recalled from an offsite location.
- High Reliability:** Storage Clouds are disk based and thus inherently more reliable than tapes. Additionally, Cloud vendors typically keep multiple redundant copies of data for availability and scalability purposes. (See AWS's [S3 Service Level Agreement](#) and [FAQs](#).)
- Unlimited Scaling and No Up-front Capital Expense:** The Cloud provides virtually unlimited capacity with no up-front capital expenditure. Consequently, users need not worry about provisioning adequate tapes or local storage to hold the required backup data. The Cloud scales seamlessly and users pay only for what they use, when they use it.
- Reduced Tape Backup and Offsite Storage Cost:** Since Cloud backup reduces or eliminates the need for tapes, this can lead to significant savings in tape backup software licensing/support and offsite tape storage costs.
- Easy Provisioning of Test and Dev Environments:** As Cloud Backups are accessible from anywhere via the Internet, they can be used to quickly clone databases to create custom test, development, or QA environments. For instance, Cloud Backups stored in Amazon S3 can be cloned to machines running in Amazon EC2 by running a simple script that is included in the Oracle-provided Amazon Machine Images (AMIs). An AMI is a virtual machine image that allows quick provisioning of a pre-installed and pre-configured Oracle database environment on Amazon EC2.

Getting Started with Cloud Backup

This section explains how to provide the means of payment to Amazon to use their storage cloud, and how to obtain and configure the OSB Cloud module with your Oracle Database.

Sign up for Amazon S3

The first step in getting started with the Oracle Secure Backup Cloud module is to sign up for Amazon S3. This can be done by visiting the Amazon S3 website (<http://aws.amazon.com/s3>). Upon successful registration, users will be provided a pair of access identifiers called the Access Key ID and the Secret Access Key.

Register for an Oracle.com or Oracle Technology Network (OTN) Account

An Oracle.com or OTN account is required to install the Oracle Secure Backup Cloud module. New accounts may be created by visiting the OTN website (<http://otn.oracle.com>).

Install the Oracle Secure Backup Cloud Module

The next step is to download the Oracle Secure Backup Cloud module install tool from OTN's [Cloud webpage](#), and run it to install and configure Cloud backups. Oracle Amazon Machine Images (AMI) on AWS's EC2 already include this install tool. Therefore, if the database being backed up is running on Amazon EC2, there is no need to download the install tool – it can be found in the `/home/oracle/scripts/osbws` directory.

The install tool can be invoked as follows (you must supply your OTN and AWS credentials):

```
$java -jar osbws_install.jar -AWSID <AWS ID> -AWSKey <AWS Secret Key>
-otnUser <OTN User ID> -otnPass <OTN Password> -walletDir <Wallet
Directory> -configFile <Cloud Backup Configuration File Name> -libDir
<Location to store Cloud Backup Module/Library> -proxyHost ww-
proxy.yourcompany.com -proxyPort <your proxy port>

Oracle Secure Backup Database Web-Service Install Tool
OTN userid is valid.
AWS credentials are valid.
Creating new registration for this S3 user.
Created new log bucket.
Registration ID: 0f0a8aac-dad0-6254-7d70-be4ac4f112c4
S3 Logging Bucket: oracle-log-jane-doe-1
Create credential oracle.security.client.connect_string1
OSB web-services wallet created in directory /orclhome/dbs/osbws_wallet.
OSB web-services initialization file /orclhome/dbs/osbwstl.ora created.
Downloading OSB Web Services Software Library.
Downloaded 13165919 bytes in 204 seconds.
Transfer rate was 64538 bytes/second.
```

```
Download complete.  
Extracted file /orclhome/lib/libosbws11.so
```

Example 1: Running the Cloud Backup Install Tool

Example 1 above shows how the tool automatically performs all the required steps to install and configure the Cloud backup module – downloading the software, creating a wallet containing the user’s AWS identifiers, and creating the Cloud backup configuration file. More details on how to run the install tool and the description of all of its arguments can be found in the install tool readme document.

Configure Recovery Manager (RMAN) Settings

This step stores the configuration information for the Cloud Backup module in the RMAN repository so that it does not need to be specified each time a backup is invoked.

```
RMAN> configure channel device type sbt parms  
'SBT_LIBRARY=/orclhome/lib/libosbws11.so  
ENV=(OSB_WS_PFILE=/orclhome/dbs/osbwst1.ora)';  
  
using target database control file instead of recovery catalog  
new RMAN configuration parameters:  
CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' PARMS  
'SBT_LIBRARY=/orclhome/lib/libosbws11.so  
ENV=(OSB_WS_PFILE=/orclhome/dbs/osbwst1.ora)';  
new RMAN configuration parameters are successfully stored
```

Example 2: Configuring RMAN

Once the RMAN configuration is completed, Cloud Backups can be performed using the same RMAN commands you usually use. This step is optional but strongly recommended.

Cataloguing and Using Cloud Backups

All Cloud backup operations will be catalogued by RMAN in the same manner as local disk or tape backups are, ensuring a seamless restore/recovery process. When a restore/recover operation is initiated, RMAN and Oracle Secure Backup Cloud module will automatically restore the required data from the Cloud - without requiring any special user intervention.

Cloud Backup Best Practices

Securing Data in the Cloud

Oracle strongly recommends encrypting your Cloud backups. Encrypting backups ensures that your data remains secure and protected against unauthorized access. Please refer to Oracle Backup and Recovery Guide to learn more about the RMAN commands that are used to configure backup encryption. Encryption can also be enabled while scheduling backups in Enterprise Manager.

Optimizing Cloud Backup Performance

As Cloud Backups are sent over the public Internet, performance is dependent on Internet network throughput – typically less than 1 MB/Sec per connection. Additionally, Cloud vendors may throttle sessions to prevent individual users from consuming disproportionate amounts of resources. According to internal tests conducted at Oracle, Amazon S3 limits an individual session’s read/write throughput to around 2-3 MB/Sec. However by using the right combination of parallelism and compression, backup speeds of up to 40-50 MB/Sec were attained; test results are summarized in Table 1, below. Some observations follow from these tests:

- Cloud Backups of your on-premise (off-cloud) databases are slower than for databases running on EC2. This is due to public Internet network bandwidth constraints.
- Compression helps overcome the network bandwidth limitations. For a database at Oracle HQ, the use of compression resulted in a 4X gain in backup speed.
- Using parallel streams (RMAN channels) also speeds up Cloud backups – particularly for on-premise databases. As can be seen in Table 1, peak performance for a database at Oracle HQ was achieved with 64 channels.

Oracle recommends the following to optimize the performance of Cloud Backups:

- Use multiple RMAN channels for higher parallelism resulting in full utilization of the network.
- Use multi-section backups. Oracle Database 11g allows multiple channels to back up a single file in parallel, increasing parallelism beyond the number of datafiles to be backed up. For example, the RMAN command to specify backup section size 1 GB is:

```
BACKUP DEVICE TYPE SBT DATABASE SECTION SIZE 1g;
```
- Use the Oracle Database 11g Advanced Compression. Oracle Database 11g Compression is significantly faster and more efficient (in terms of CPU overhead) than pre-11g compression.
- Consider making full database backups once a week and performing incremental backups during the weekdays. This will result in faster backups and may help save significant amount of network bandwidth. Use the RMAN Block Change Tracking feature to optimize the performance of your daily incremental backups.

Test Environment	Uncompressed Backup Speed (Network Throughput)	Compressed Backup Speed	Full DB Backup Time (250 GB)	Incremental Backup Time (10% delta)
DB at Oracle HQ (8 x 2 GHz CPU, 16 GB RAM)	10 MBPS (64 RMAN Channels)	40 MBPS (64 RMAN Channels)	2-6 Hours	30 Minutes – 1 Hour
DB within Amazon Cloud (Extra Large EC2 Instance)	35 MBPS (16 RMAN Channels)	50 MBPS + Constrained by CPU (32 RMAN Channels)	2 Hours	< 20 Minutes

Table 1: Cloud Backup Performance

Conclusion

The Oracle Database Cloud Module allows customers to use Amazon’s Simple Storage Service (S3) as their offsite backup storage destination. Compared to traditional tape-based offsite storage, Cloud backups are more accessible, faster to restore under most circumstances, and more reliable, while eliminating the overheads associated with maintaining off-site backup operations.

Cloud backups are also the optimal protection for databases running within the compute Cloud.



Cloud-Stored Offsite Database Backups
May 2010

Author: Cris Pedregal

Contributing Authors: Bill Hodak, Muthu
Olagappan

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2010, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0110