

Oracle Database Backup in the Cloud

*An Oracle White Paper
September 2008*

Oracle Database Backup in the Cloud

Introduction	2
Cloud backup Overview.....	2
Introducing Oracle Secure Backup Cloud Module.....	3
Complete Data Security and Privacy with Built-In Encryption	4
Intelligent, Network-Optimized Backup	4
Database Version Support	4
Cloud Backup Benefits.....	5
Amazon Simple Storage Service (S3)	5
Getting started with Cloud Backup	6
Sign up for Amazon S3 Service	6
Register for an Oracle.com or Oracle Technology Network (OTN) Account	6
Install the Oracle Secure Backup Cloud Module	6
Configure Recovery Manager (RMAN) Settings.....	7
Monitoring Cloud Backups	8
Cloud Backup Best Practices	8
Securing Data in the Cloud	8
Optimizing Cloud Backup Performance	8
Conclusion.....	10

Oracle Database Backup in the Cloud

INTRODUCTION

Cloud Computing has emerged as the new industry buzzword. While there are several definitions of what constitutes a Cloud and what the term Cloud Computing really means, simply put, Cloud Computing allows users to tap into a virtually unlimited pool of computing resources over the Internet (the Cloud). Unlike traditional IT, Cloud users have little insight or control over the underlying infrastructure, and they must interact with the Cloud via an API provided by the Cloud vendors. The most exciting part of the Cloud, however, is its elasticity and utility-style pricing. Users can provision resources dynamically in a self-service manner and pay only for what they use – much like how we use utilities like electricity or water.

While the Cloud has created a loyal set of followers – mostly among web startups and developers – enterprises have found it somewhat difficult to fully utilize due to a variety of reasons, including unavailability of enterprise software, data security and privacy issues, etc. Oracle is addressing this challenge by introducing products and services that will help make the Cloud enterprise-ready. Using the new Cloud offerings from Oracle, enterprises will be able to leverage the Cloud to meet transient and unanticipated demands by treating the Cloud as an extension to their on-premises data center.

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 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 BACKUP OVERVIEW

Enterprises have traditionally relied on tapes for offsite backups. Tapes are transported to an offsite storage location and stored there for a pre-determined period, following which they are returned and recycled.

Introducing Oracle Secure Backup Cloud Module

New disk economics and the compelling price points offered by Storage Cloud vendors present new opportunities to make offsite backups more accessible and reliable. Responding to this, Oracle is pleased to introduce the capability to backup Oracle Database to the Cloud. Using the newly introduced Oracle Secure Backup Cloud module, it is now possible to send local disk backups directly to the Storage Cloud for offsite storage purposes. The Cloud backup module is a part of the Oracle Secure Backup product family. 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.

The Oracle Secure Backup Cloud module can also be used to stream backups directly to the Cloud. This is particularly useful when the database is also running in the Cloud, using services such as Amazon Elastic Compute Cloud (EC2).

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, the Oracle Secure Backup Cloud module can be used to back up all supported versions of Oracle Database i.e. Oracle Database 9i Release 2 or higher. In addition, database administrators can continue to use their existing backup tools – Enterprise Manager, scripts, etc. – to perform Cloud backups.

Oracle Secure Backup Cloud module is currently available for Linux and Windows (32-bit only) platforms.

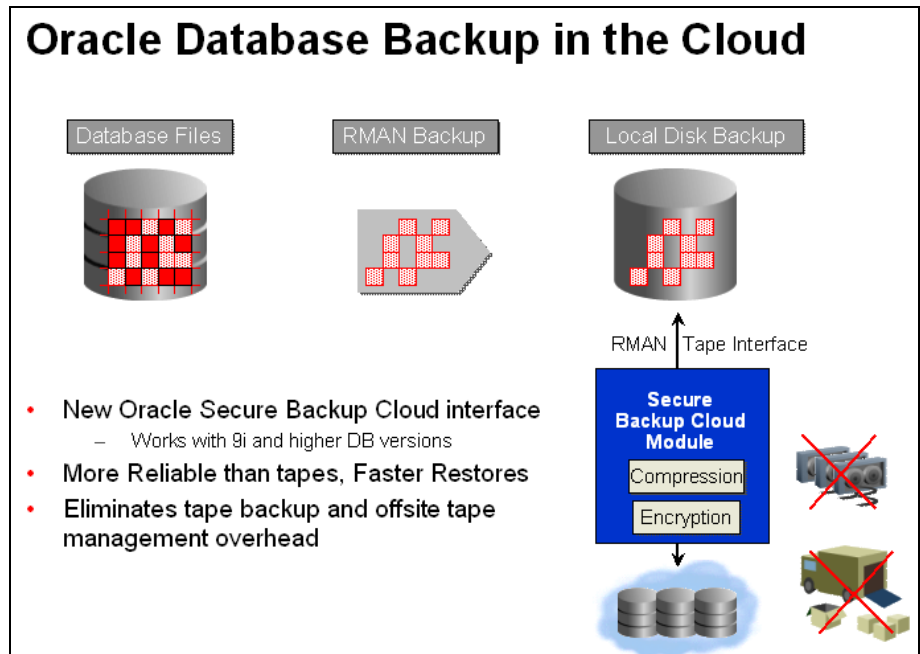


Fig 1: Oracle Database backup in the Cloud

Complete Data Security and Privacy with Built-In Encryption

Oracle Secure Backup provides the ability to encrypt backups to ensure data security and privacy. 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 Secure Backup's ability to encrypt data at its source further mitigates risk of theft or unauthorized access because the backup data remains encrypted both in-transit and at rest in the Cloud.

Intelligent, Network-Optimized Backup

Network bandwidth consumption and the time required to complete backups can be another major concern while backing up databases to the Cloud, over the public Internet. The Oracle Secure Backup Cloud module features a rich set of capabilities to intelligently minimize the volume of backup data and also to compress backups. Thanks to its tight integration with the Oracle Database engine, Oracle Secure Backup can identify unused space (blocks) within the database and skip backing it up. Since database administrators typically allocate more space than they actually use, to provide headroom for future growth, this can lead to a substantial reduction in both backup size and time. Additionally, starting with Oracle Database 11g, Oracle Secure Backup can also eliminate committed undo data from backup by identifying and backing up only the "active" undo data (i.e. undo data related to transactions in progress at the time of backup). Database undo sizes typically range from 10 to 100s of GB and this optimization can also lead to significant improvements in backup performance.

It is important to note that both unused space and undo data backup optimization are unique features of Oracle Secure Backup and are not available with any other backup product.

With Oracle Secure Backup Cloud module, users can also utilize RMAN's rich compression capabilities. This includes the Oracle Database 11g Fast Compressed Backup feature, which normally requires the Advanced Compression option. However, Oracle Secure Backup Cloud module users can use this functionality for no additional cost. Compression typically results in a 50% reduction in backup size. 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 Program can or may be used to backup the following supported versions of Oracle Database: Oracle Database 9i Release 2 or higher.

Cloud Backup Benefits

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

- **Continuous Accessibility, Faster Restores:** 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.
- **Better Reliability:** Storage Clouds are disk based and thus inherently more reliable than tapes. Additionally, the Cloud vendors typically keep multiple redundant copies of data for availability and scalability purposes.
- **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.
- **Quick 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, dev, 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.

Amazon Simple Storage Service (S3)

Amazon Simple Storage Service (<http://aws.amazon.com/s3>) is part of the Amazon Web Services (AWS) offering. Amazon S3 provides a simple web-services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web. It provides users access to the same highly scalable, reliable, fast, inexpensive data storage infrastructure that Amazon uses to run its own global network of web sites.

Amazon S3 is the first Cloud vendor that Oracle has partnered with to enable database backup in the Cloud. This functionality may be extended to more Cloud platforms/vendors in the future.

GETTING STARTED WITH CLOUD BACKUP

Sign up for Amazon S3 Service

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*. These access identifiers can be retrieved from the Amazon S3 web site by hovering the mouse over the “Your Web Services Account” button on the upper right hand side corner of the page and then clicking on the “AWS Access Identifiers” link.

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>) and clicking on the “Sign In/Register for Account” link at the top of the page.

Install the Oracle Secure Backup Cloud Module

The next step is to download the Oracle Secure Backup Cloud module install tool from OTN, and run it to install and configure Cloud backups. The install tool is pre-packaged with the AMIs provided by Oracle to ease Oracle software deployment on Amazon EC2. 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:

```
$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 www-
proxy.smallcompany.com

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/osbwst1.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

As can be seen from the above sample output, the install tool automatically performs all the required steps to install and configure the Cloud backup module – including 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 optional, but highly recommended, 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 above step is completed, the system is fully configured for Cloud backups. Cloud Backups can now be performed using the same RMAN commands you usually use.

Similarly, all Cloud backup operations will be catalogued by RMAN in the same manner as local disk or tape backups are, thereby ensuring a completely transparent and seamless restore/recovery process. In another words, 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.

Monitoring Cloud Backups

Oracle is making available a free web portal – Oracle Cloud Control – to provide basic monitoring and reporting for Cloud backups. This portal provides information such as total cloud storage used for storing backups, observed network performance for backup and restore operations, and other such details. It can be accessed from OTN using the OTN username/password provided at the time of running the install tool.

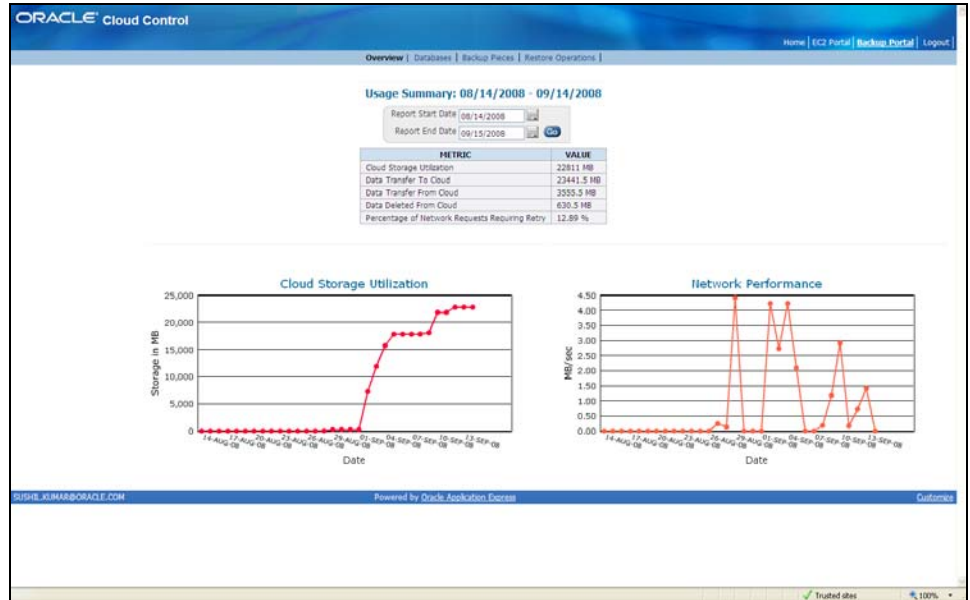


Figure 2: Oracle Clouds Backup Reporting Portal

CLOUD BACKUP BEST PRACTICES

Securing Data in the Cloud

Oracle strongly recommends encrypting your Cloud backups. As stated earlier, encrypting backups ensures that your data remains secure and protected against any kind of unauthorized access. Please refer to Oracle Backup and Recovery Guide (http://download.oracle.com/docs/cd/B28359_01/backup.111/b28270/rmcconf_a.htm#BRADV89467) 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 done over the public Internet, their performance is highly dependent on Internet network throughput – typically less than 1 MB/Sec per connection. Additionally, Cloud vendors may also 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. The following table summarizes the details of the tests and observed results.

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 or Higher Constrained by CPU (32 RMAN Channels)	2 Hours	< 20 Minutes

Table 1: Cloud Backup Performance

The following conclusions can be derived from the above results:

- Cloud Backups are going to be slower for your on-premise databases than for databases running on EC2. This is due to public Internet network bandwidth constraints.
- Compression can be a very effective tool to overcome the network bandwidth limitations. For a database at Oracle HQ, the use of compression resulted in a 4x gain in backup speed.
- A large number of parallel concurrent streams (RMAN channels) also helps in speeding up Cloud backups – particularly for on-premise databases. As can be seen above, peak performance for a database at Oracle HQ was achieved with 64 channels.

Based on the above test results, Oracle recommends the following to optimize the performance of Cloud Backups:

- Use multiple RMAN channels. Prior to Oracle Database 11g, each datafile could only be backed up by one RMAN channel. Consequently, the effective degree of parallelism is limited by the number datafiles in the database. Oracle Database 11g however has intra-datafile parallelism capability that allows multiple channels to backup a single file in parallel.

- Compress your backups and use the new Oracle Database 11g Fast Compressed Backup feature. Compared to pre-11g backup compression, Oracle Database 11g Fast Backup Compression is significantly faster and more efficient (in terms of CPU overhead).
- Consider making full database backups once a week and perform 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.

CONCLUSION

Exciting developments taking place in the area of Cloud Computing present new opportunities for enterprises. By introducing the capability to backup databases in the Cloud in a secure and efficient manner, Oracle is leading the effort to make the Cloud relevant to enterprises. Backups in the Cloud provide enterprises with a next-generation offsite backup storage solution that is always accessible, faster to restore under most circumstances, and more reliable. Continuous accessibility of Cloud backups opens up new possibilities that have not been possible with tape based offsite storage. With its unlimited capacity and extremely attractive consumption-based pricing, the Cloud can help enterprises avoid up-front capital expenditure and become more agile, efficient, and competitive - with Oracle helping pave the way.



Oracle Database Backup in the Cloud
September 2008
Authors: Sushil Kumar

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

Copyright © 2007, Oracle. 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 is a registered trademark of Oracle Corporation and/or its affiliates.

Other names may be trademarks of their respective owners.