Oracle Total Recall/Flashback Data Archive

An Oracle White Paper
June 2007

ORACLE

INTRODUCTION

Organizations are faced with increasing data retention and change control requirements. The major phenomenon fueling this trend is regulatory oversight and compliance. Regulations such as Sarbanes-Oxley, HIPAA and Basel-II enforce strict change control and tracking requirements for customer data. These regulations affect companies both big and small, in all industry segments. Most of these regulations specify long retention periods, typically 5 or more years, during which the history should be readily accessible. Stop gap or short-term approaches to managing history data do not work. Companies need a robust, long-term history management system. Organizations are further realizing the business value that can be extracted from the historical or temporal dimensions of data. Valuable customer information can be mined from the historical data store. Failure to maintain appropriate tracking and retention is expensive. Non-compliance could lead to fines, loss of investor confidence and reputation.

A variety of approaches are currently employed to create and manage historical transaction data. Unfortunately all of them come with significant limitations and force organizations to make costly compromises. Flashback Data Archive is a new feature in Oracle Database 11g that can automatically track and maintain historical changes to all Oracle data in highly application transparent, secure and efficient manner. Part of the Oracle Total Recall Option, Flashback Data Archive overcomes the limitations of the current approaches to historical data management and provides enterprises with a quick, centralized and extremely efficient solution to meet their historical data management needs.

HISTORICAL DATA MANAGEMENT: CURRENT APPROACHES

A popular approach to implementing a historical data management system is at the application level. The logic for history management and archiving is embedded within the business logic. A key advantage of this approach is that applications understand the business context and can provide easy grouping of different business components that need to be tracked. However, this comes with a cost: increased application complexity. Data access has to be through defined application interfaces to ensure accurate and complete history tracking. Direct updates to the data at the database level could introduce data integrity issues with the history and

more importantly, could invalidate the whole history. In an organization where there are typically 100s of applications, with a mix of third party and homegrown applications, implementing the historical data management system in every application is extremely difficult. This problem is further accentuated by application upgrades, changing retention requirements and a lack of centralized management interface.

A second approach is to use database triggers to maintain data history. This approach solves many of the problems associated with application-managed data history. It, however, suffers from limitations such as a lack of centralized management interface, which leads to increased maintenance overhead. Moreover, since triggers are executed with every transaction in the foreground, user processes and applications could suffer from degraded performance.

Finally, there are third-party solutions that mine redo logs to create data history. Since history is stored in a separate database, this approach cannot provide seamless access to the OLTP and history data. Additionally, these might require extra training for users due to different interfaces and are often expensive and difficult to build. None of these approaches provide a high-performance, secure and seamless history data management system.

FLASHBACK DATA ARCHIVE

The Flashback Archive Data feature in Oracle Database 11g – part of the Oracle Total Recall option - overcomes most of the above limitations. It automatically tracks every single change made to the data stored inside the database and maintains a secure, efficient and easily accessible archive of historical data. The captured historical data can be retained for as long as the business demands and is easily accessible using standard SQL constructs. Historical data tracking can be enabled on both existing and new tables instantaneously and more importantly, in a completely application transparent manner.

Implemented natively inside the database, Flashback Data Archive presents a high-performance, storage optimized solution with a centralized management interface for satisfying data retention and change control requirements for organizations. The primary advantages of using Flashback Data Archive for historical data tracking include:

1. Application Transparent – Enabling historical data capture on one or more tables can be done instantaneously with no or minimal application changes. Customers can therefore use this feature to capture historical data for both packaged as well as home grown applications.

2. Seamless Access – Historical data can be easily accessed using familiar Flashback SQL constructs. Flashback Data Archive includes support for Flashback Queries. Applications can seamlessly query the history of table data, as it existed in different points in time. No special snapshots need to be taken to take advantage of this feature.

3. Security – historical data, once generated, is immutable to all users. This is enabled out-of-the-box and no special or extra setup is required. Access to the internal history tables is restricted to reads only. No DML operations are allowed to users, including administrators. Applications need not query the internal history tables directly as seamless access is provided through the Flashback Query mechanism.

4. Minimal performance overhead – Regular user transactions will see negligible impact. Flashback Data Archive employs a lightweight mechanism to mark DML operations on tracked tables for archiving. The actual history generation and archiving is done asynchronously through a background process as explained later.

5. Storage Optimized – The history data is internally partitioned and highly compressed to reduce the storage footprint. Flashback Data Archive employs a highly efficient compression scheme to compress the internal history tables. In addition, it automatically partitions the internal history tables based on a range-partitioning scheme. Both compression and partitioning in flashback data archive are managed automatically and require no special administration.

6. Centralized Management – Flashback Data Archive provides a centralized and policy-based management interface to automate a number of ongoing administrative tasks. With Flashback Data Archive, you can easily group tables and set a common retention policy. New tables will automatically inherit the retention parameter from the flashback data archive it is a part of. Oracle will automatically purge aged-out history data for all tracked tables based on the specified retention. This frees up the administrator from the repetitive management of history data and avoids costly errors associated with manual maintenance such as purging wrong history.

ORACLE FLASHBACK TECHNOLOGY

Oracle's Flashback technology provides a set of functionality to access data as of a time in the past and recover from human errors. Flashback technology is unique to the Oracle Database and supports recovery at any level of granularity including the row, transaction, table, and database wide. Using flashback features you can query past versions of data as well as perform change analysis and self-service repair to recover from logical corruption while the database is online. Flashback technology includes the following features:

- **Flashback Query** allows the user to query data at some point-in-time in the past to reconstruct lost data that may have been deleted or changed by accident

- **Flashback Version Query** provides a mechanism to view changes made to the database over time at the row level.

- **Flashback Transaction Query** provides a mechanism to view changes made to the database at the transaction level.
- **Flashback Database** is a new strategy for doing point-in-time recovery. It quickly rewinds an Oracle database to a previous time to correct any problems caused by logical data corruption or user error.
- **Flashback Table** provides the ability to recover a table, or a set of tables, to a specified point in time in the past very quickly and easily.
- **Flashback Drop** provides a safety net when dropping objects as you can very quickly and easily undrop a table and its dependent objects.

Flashback Data Archive is the newest addition to the Flashback family. In Oracle Database 11g Release 1, Flashback Data Archive provides support for querying the history data using the familiar Flashback Query SQL construct 'AS OF'. An important distinction that distinguishes Flashback Data Archive from the other Flashback features is that history data can be accessed for any specified duration, far and beyond the amount of retained undo or flashback log data. Historical data can be retained for any length of time as required for compliance or other business reasons. With Flashback Data Archive, history data is always available and readily queryable.

ARCHITECTURE

Oracle implements a multi-versioning mechanism that ensures read-consistency while maintaining high degree of concurrency. When DML operations such as insert, update or delete happen on data, Oracle writes the data into an undo tablespace that is used not only for transaction rollbacks but also guaranteeing read consistency in a concurrent environment. Flashback Data Archive works with the undo data as with many other Flashback features. As stated above, however, the historical data retain in the flashback data archive is no limited by the size of undo tablespace. History generation is implemented in a non-intrusive manner through a

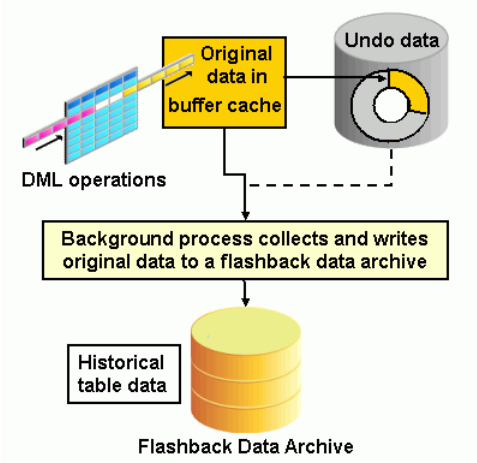

new background process called 'fbda'. After a table has been enabled for history

tracking with Flashback Data Archive, all transactions and the corresponding undo records are marked for archival. In order to guarantee that every transaction is archived, the undo records are not recycled until the history is generated and stored in the database. The process sleeps and wakes up at system-managed intervals and processes the undo data marked for archival. After 'fbda' generates the history, the transactions and undo records are candidates for recycling. This asynchronous 'fbda' process ensures negligible impact on the overall transaction or foreground performance.

The 'fbda' process intelligently adjusts its sleep interval based on the system undo generation rate. As transaction activity increases, 'fbda' automatically reduces the sleep interval from the default sleep time of 5 minutes. For better performance, 'fbda' process also adjusts its sleep time in order to maximize undo data reads from the buffer cache. Flashback Data Archive also uses an internal partitioning scheme for all the historical data for better performance. It also employs table compression to reduce the storage footprint of the historical data that could easily grow into hundreds of terabytes for longer retention times.

UNDERSTANDING FLASHBACK DATA ARCHIVE

A Flashback Data Archive is a logical container for managing historical information for related tables. It is a new data dictionary object in Oracle Database 11g that defines archive storage and data purging policies A Flashback Data Archive can span multiple tablespaces. Administrators can define the amount of space a flashback data archive can use in each tablespace using the 'QUOTA' parameter. Multiple flashback data archives can be created as needed to implement different archiving policies. Each flashback data archive contains a 'RETENTION' parameter that specified the duration for retaining the historical changes. Flashback Data Archive guarantees that historical data will be retained for the duration specified by 'RETENTION' and automatically purges out aged historical data. Additionally, administrators, with the necessary privileges, can purge historical data in an ad-hoc manner.

Flashback Data Archive creates an internal history table for every tracked table. The internal history table is a replica of the tracked table with additional timestamp columns. When one or more columns in the tracked table are updated, a new row is inserted into the history table that is the before-image of the row before the transaction. It is important to note that UPDATE and DELETE operations generate a new record in the history table. Flashback Data Archive does not create a new history record for INSERT operations. The internal history table is partitioned for better performance. No modifications are allowed to internal partitions. The internal history tables are compressed to reduce the disk space requirements. Applications or users need not access the internal history tables directly. The 'AS OF' SQL construct can be used to seamlessly query the historical data.

HISTORY TRACKING USING FLASHBACK DATA ARCHIVE

Flashback Data Archive Requirements

1. Tablespaces managed with automatic segment space management (ASSM) are required for using with Flashback Data Archive
2. Automatic Undo Management must be enabled

Step-by-Step Approach

There are 3 steps to enabling history tracking on tables using Flashback Data Archive.

1. Create a new tablespace with automatic segment space management for storing historical data (Optional)
2. Create a flashback data archive

```
CREATE FLASHBACK ARCHIVE fda1
TABLESPACE tbs1
RETENTION 5 YEAR;
```

This creates a new flashback data archive called 'fda1' on tablespace tbs1 with a retention policy of 5 years. Since 'QUOTA' is not specified, the default setting of Unlimited is used. At this point, we have created a logical container or placeholder for historical data.

- FLASHBACK ARCHIVE ADMINISTER system privilege is required for creating a new flashback data archive
- The following static data dictionary views are available
  - DBA/USER_FLASHBACK_ARCHIVE – Displays information about Flashback Data Archives
  - DBA/USER_FLASHBACK_ARCHIVE_TS – Displays tablespaces and the mapping to Flashback Data Archives

3. Enable history data archiving desired tables

```
ALTER TABLE EMPLOYEES FLASHBACK ARCHIVE fda1;
```

Table EMPLOYEES now has been enabled for historical data tracking and its history data will be stored in flashback data archive 'fda1'. When the first DML operation happens, the internal history table is created in 'fda1'.

- FLASHBACK ARCHIVE object privilege is required to enable historical data tracking
- The following static data dictionary views are available
  - DBA/USER_FLASHBACK_ARCHIVE_TABLES - Displays information about tables that are enabled for flashback archiving.

QUERYING HISTORY DATA

Flashback Data Archive provides seamless access to the historical data using the 'AS OF' SQL construct. You can query for the state of any row in the tracked table going as far back as your specified retention period.

The following is an example for querying the salary details for employee with id=193 on June 1, 2007.

```
SELECT last_name, first_name, salary

FROM EMPLOYEES AS OF TIMESTAMP TO_TIMESTAMP ('2007-06-
01 00:00:00', 'YYYY-MM-DD HH24:MI:SS')

WHERE employee_id=193;
```

MANAGING FLASHBACK DATA ARCHIVES

Flashback Data Archive guarantees historical data capture for all transactional changes to tracked tables. Flashback Data Archive therefore automatically disallows any operation that would invalidate history. DDL operations that are disallowed on tracked tables include dropping, truncating or renaming tables, dropping or modifying columns. Partition or subpartition operations on the internal history tables are also disallowed. It is however possible to add new columns without disrupting historical data capture.

New Privileges for Flashback Data Archive

- FLASHBACK ARCHIVE ADMINISTER – A new system privilege managing flashback data archives. Allowed operations include:
  - Create new flashback data archive
  - Modify existing flashback data archive
    - Make a specific flashback data archive as default
    - Set RETENTION time
    - Ad-hoc purge of history data
    - Add/remove tablespaces
  - Disable tracking for tracked tables
  - Drop flashback data archive
- FLASHBACK ARCHIVE – A new object privilege for enabling tracking on tables on the flashback data archives

RECOMMENDED BEST PRACTICES

In order to optimize performance and space consumption, Flashback Data Archive does not automatically replicate base table indexes on the internal history tables. Users are recommended to query the DBA/USER_FLASHBACK_ARCHIVE_TABLES view for the name of the internal table using and create the desired indexes on this internal history table. It is important to note that these indexes will have a performance impact on the Flashback Data Archive process as it will have to maintain the indexes in addition to collecting and archiving the transactional changes.

Secondly, it is recommended to use dedicated tablespaces for flashback data archives. The QUOTA parameter should be left at the default value of Unlimited. This makes space planning and maintenance of flashback data archives easier. Also, group tables based on desired retention period and create a new flashback data archive for each retention period.

Finally, access to the system privilege 'FLASHBACK ARCHIVE ADMINISTER' should be restricted. It is recommended to separate user access to the system privilege 'FLASHBACK ARCHIVE ADMINISTER' from the object privilege 'FLASHBACK ARCHIVE'.

CONCLUSION

Flashback Data Archive provides a secure, efficient and application transparent solution for generating and managing historical data. It provides a centralized and integrated interface to manage and retain data history. An automated, policy-based management greatly simplifies administration. With Flashback Data Archive, you can easily implement historical data tracking to comply with new regulations or adapt to changing business needs.

White Paper Title
June 2007
Author: RAVI RAJAMANI

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