

**19<sup>c</sup>** ORACLE<sup>®</sup>  
Database

# Oracle Flashback Data Archive (FDA)

ORACLE WHITE PAPER | FEBRUARY 2019



ORACLE<sup>®</sup>



## Table of Contents

Disclaimer	1
<b>Introduction</b>	2
<b>Tracking/Viewing Changes is Complicated</b>	3
<b>Enabling Flashback Data Archive</b>	4
<b>Flashback Data Archive Usage Example</b>	8
<b>Best Practices and More Information</b>	11
<b>Conclusion</b>	11

## Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.



## Introduction

Recreating history, after the fact, is expensive and often impossible.

Yet every day IT organizations, and end-users, are faced with this data dilemma. They know that their data has been modified since it was originally created. They could easily query the latest value(s) from the database, but what they really need to know is the value from 2 days ago, before the data was modified.

It is not hard to imagine the everyday use cases -- U.S. government regulations such as the Sarbanes-Oxley Act and the Health Insurance Portability and Accountability Act (HIPAA) or international agreements such as Basel II are driving the need to track changes to data over long periods of time. Business and service organizations are increasingly held accountable for maintaining a secure history of data on which key business decisions, that often have both legal and business bottom-line ramifications, are based -- having a reliable history of change to your data is business critical.

The Flashback Data Archive (FDA) feature of Oracle Database provides a mechanism for tracking changes to production databases that is secure, efficient, easy to use and application transparent. FDA allows organizations to configure historical data capture in a matter of minutes, providing a centralized and seamlessly queryable historical data store that makes the most efficient use of all your resources – CPU, storage or administration time.

This document will discuss how FDA works, how FDA is configured and provides usage examples of how FDA can be used to maintain an organizations historic data.

## Tracking/Viewing Changes is Complicated

History is complicated.

- Extensive application logic is needed to track history
- Application logic must be maintained with every application change

Viewing history is complicated.

- Querying and reporting history data is hard, as schemas evolve
- Viewing data across multiple history tables is even harder

Prior to Flashback Data Archive, one of the approaches was to capture and manage historical data at the application level. The perceived advantage of this approach is that this approach enabled the best understanding of the data relationships and dependencies at the application level – yet the possibility of direct data access (outside the application) could cause change tracking issues.

### Approach: Application Logic

- **Problem:** Direct database updates bypass history capture

The possibility that someone changes the data directly, outside the application, is of serious concern. For one, it increases application complexity as the historical data management logic needs to be embedded into every single application. Similarly, it may be difficult to implement retention policies, as there is no centralized management of the historical data. Finally, organizations would need to be vigilant and always keep in mind that data is not always accessed through applications and if someone were to make a change directly at the database level, or using any other non-application interface, no history will be captured.

Problems with some of the additional approaches before FDA included:

### Approach: Database Triggers


- **Problem:** Performance impact

Some organizations have captured history at the database level using triggers. While it does resolve some of the issues of application level history capture, the trigger-based mechanism could lead to significant performance and maintenance overhead.

### Approach: Log Mining

- **Problem:** Complex & error prone

Log mining techniques could be employed, provided the nuances and complexity is handled correctly, which can be complex as well.



Oracle Flashback Technology is a group of Oracle Database features that let you view past states of database objects or to return database objects to a previous state without using point-in-time media recovery.

Oracle Flashback features uses the Automatic Undo Management (AUM) system to obtain metadata and historical data for transactions. They rely on **undo data**, which are records of the effects of individual transactions. The background process FBDA will start to automatically start, manage, and coordinate the task of archiving transactions and maintaining history tables. Archiving of individual transaction is either done by FBDA process, slave processes coordinated by FBDA, or inline by foreground processes doing the DMLs.

With flashback features, you can:

- Perform queries that return past data
- Perform queries that return metadata that shows a detailed history of changes to the database
- Recover tables or rows to a previous point in time
- Automatically track and archive transactional data changes
- Roll back a transaction and its dependent transactions while the database remains online

Flashback Data Archive provides the ability to track and store transactional changes to a table over its lifetime. Flashback Data Archive is useful for compliance with record stage policies and audit reports. The Flashback Data Archive feature in Oracle Database (previously known as the Total Recall feature in Oracle Database 11g), provides a mechanism for tracking changes to production databases that is secure, efficient, easy to use and application transparent.

Flashback Data Archive consists of one or more tablespaces or parts thereof. You can have multiple Flashback Data Archives. Flashback Data Archive is configured with retention time. Data archived in the Flashback Data Archive is retained for the retention time specified when the Flashback Data Archive was created. FDA is supported for multitenant for both local and shared undo configurations. Flashback archives are created at per PDB level just as user tables are created and enabled for flashback archive at a PDB level. Retention property is associated with a Flashback Archive and will be honored by the PDB.

## Enabling Flashback Data Archive

This section provides an overview of how FDA is enabled, for more details please see [here](#).

## FDA Restrictions and Best Practices

Data archived in the Flashback Data Archive is retained for the retention time specified when the Flashback Data Archive was created.

By default, Flashback Data Archive is not enabled for any tables. You can enable Flashback Data Archive for a table if all of these conditions are true:

- You have the FLASHBACK ARCHIVE object privilege on the Flashback Data Archive to use for that table.
- The table is not a nested table, temporary table, external table, materialized view, Advanced Query (AQ) table, or non-table object.
- The table contains neither LONG nor nested columns.

The table does not use any of these Flashback Data Archive reserved words as column names:

STARTSCN  
ENDSCN  
RID  
XID  
OP  
OPERATION

After Flashback Data Archive is enabled for a table, you can disable it only if you either have the FLASHBACK ARCHIVE ADMINISTER system privilege or you are logged on as SYSDBA.

When choosing a Flashback Data Archive for a specific table, consider the data retention requirements for the table and the retention times of the Flashback Data Archives on which you have the FLASHBACK ARCHIVE object privilege.

## FDA Creation

Create a Flashback Data Archive with the CREATE FLASHBACK ARCHIVE statement, specifying:

- Name of the Flashback Data Archive
- Name of the first tablespace of the Flashback Data Archive
- (Optional) Maximum amount of space that the Flashback Data Archive can use in the first tablespace

The default is unlimited. Unless your space quota on the first tablespace is also unlimited, you must specify this value; otherwise, error ORA-55621 occurs.

- Retention time (number of days that Flashback Data Archive data for the table is guaranteed to be stored)

- (Optional) Whether to optimize the storage of data in the history tables maintained in the Flashback Data Archive, using [NO] OPTIMIZE DATA.

The default is NO OPTIMIZE DATA.

If you are logged on as SYSDBA, you can also specify that this is the default Flashback Data Archive for the system. If you omit this option, you can still make this Flashback Data Archive the default later.

Oracle recommends that all users who must use Flashback Data Archive have unlimited quota on the Flashback Data Archive tablespace; however, if this is not the case, you must grant sufficient quota on that tablespace to those users. Best practice is to create a separate flashback data archive for each application module you archive.

#### **Examples:**

Create a default Flashback Data Archive named fla1 that uses up to 10 G of tablespace tbs1, whose data are retained for one year:

```
CREATE FLASHBACK ARCHIVE DEFAULT fla1 TABLESPACE tbs1
QUOTA 10G RETENTION 1 YEAR;
```

Create a Flashback Data Archive named fla2 that uses tablespace tbs2, whose data are retained for two years:

```
CREATE FLASHBACK ARCHIVE fla2 TABLESPACE tbs2 RETENTION 2 YEAR;
```

#### **FDA Enabling/Disabling**

By default, Flashback Data Archive is disabled for any table. You can enable Flashback Data Archive for a table if you have the FLASHBACK ARCHIVE object privilege on the Flashback Data Archive to use for that table.

To enable Flashback Data Archive for a table, include the FLASHBACK ARCHIVE clause in either the CREATE TABLE or ALTER TABLE statement. In the FLASHBACK ARCHIVE clause, you can specify the Flashback Data Archive where the historical data for the table are stored. The default is the default Flashback Data Archive for the system. If you specify a nonexistent Flashback Data Archive, an error occurs.

If you enable Flashback Data Archive for a table, but AUM is disabled, error ORA-55614 occurs when you try to modify the table.

If a table has Flashback Data Archive enabled, and you try to enable it again with a different Flashback Data Archive, an error occurs.

After Flashback Data Archive is enabled for a table, you can disable it only if you either have the FLASHBACK ARCHIVE ADMINISTER system privilege or you are logged on as SYSDBA. To disable Flashback Data Archive for a table, specify NO FLASHBACK ARCHIVE in the ALTER TABLE



statement. (It is unnecessary to specify `NO FLASHBACK ARCHIVE` in the `CREATE TABLE` statement, because that is the default.)

After enabling Flashback Data Archive on a table, Oracle recommends waiting at least 20 seconds before inserting data into the table, and waiting up to 5 minutes before using Flashback Query on the table.

### Examples

Create table `employee` and store the historical data in the default Flashback Data Archive:

```
CREATE TABLE employee (EMPNO NUMBER(4) NOT NULL, ENAME VARCHAR2(10),  
JOB VARCHAR2(9), MGR NUMBER(4)) FLASHBACK ARCHIVE;
```

Create table `employee` and store the historical data in the Flashback Data Archive `fla1`:

```
CREATE TABLE employee (EMPNO NUMBER(4) NOT NULL, ENAME VARCHAR2(10),  
JOB VARCHAR2(9), MGR NUMBER(4)) FLASHBACK ARCHIVE fla1;
```

Enable Flashback Data Archive for the table `employee` and store the historical data in the default Flashback Data Archive:

```
ALTER TABLE employee FLASHBACK ARCHIVE;
```

Enable Flashback Data Archive for the table `employee` and store the historical data in the Flashback Data Archive `fla1`:

```
ALTER TABLE employee FLASHBACK ARCHIVE fla1;
```

Disable Flashback Data Archive for the table `employee`:

```
ALTER TABLE employee NO FLASHBACK ARCHIVE;
```

### FDA DDL Support

Flashback Data Archive supports only these DDL statements:

- `ALTER TABLE` statement that does any of the following:
  - Adds, drops, renames, or modifies a columns
  - Adds, drops, or renames a constraint
  - Drops or truncates a partition or subpartition operation
- `TRUNCATE TABLE` statement
- `RENAME` statement that renames a table

Flashback Data Archive does not support DDL statements that move, split, merge, or coalesce partitions or subpartitions, move tables, or convert `LONG` columns to `LOB` columns.



For example, the following DDL statements cause error ORA-55610 when used on a table enabled for Flashback Data Archive:

- ALTER TABLE statement that includes an UPGRADE TABLE clause, with or without an INCLUDING DATA clause
- ALTER TABLE statement that moves or exchanges a partition or subpartition operation
- DROP TABLE statement

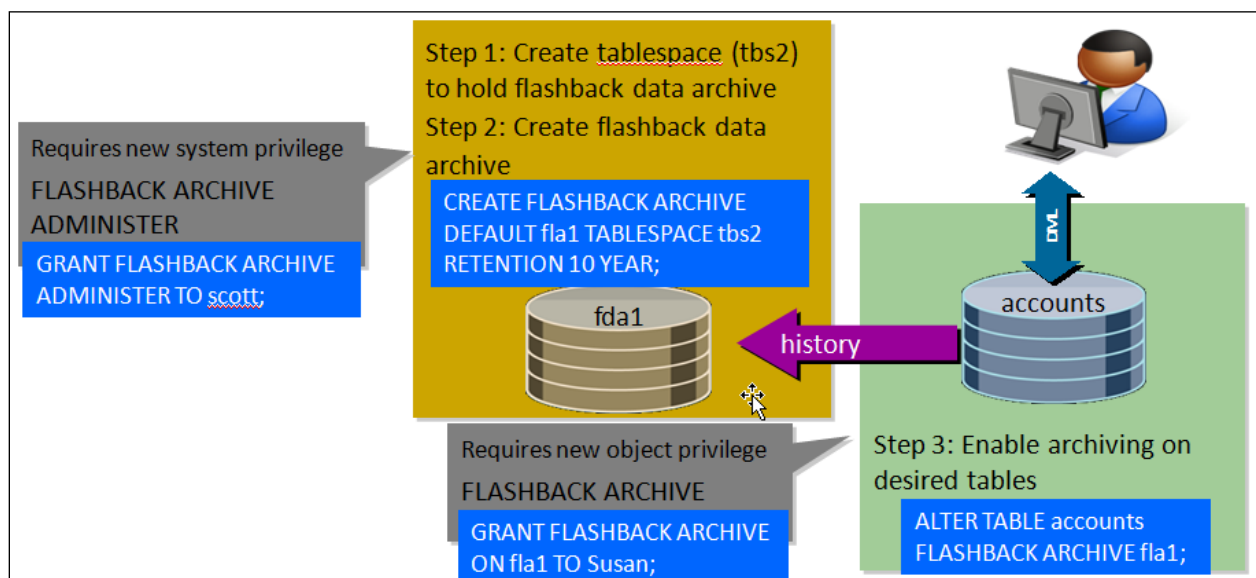
If you must use unsupported DDL statements on a table enabled for Flashback Data Archive, use the DBMS\_FLASHBACK\_ARCHIVE.DISASSOCIATE\_FBA procedure to disassociate the base table from its Flashback Data Archive. To reassociate the Flashback Data Archive with the base table afterward, use the DBMS\_FLASHBACK\_ARCHIVE.REASSOCIATE\_FBA procedure. Also, to drop a table enabled for Flashback Data Archive, you must first disable Flashback Data Archive on the table by using the ALTER TABLE ... NO FLASHBACK ARCHIVE clause.

## Flashback Data Archive Usage Example

This section provides a simple example of how FDA can be used, the example is focused the use of FDA and assumes the table schema has already been created.

Flashback Data Archive requires very little setup and administration. Two new privileges have been introduced, the Flashback Archive Administer and Flashback Archive privilege. The Administer creates a tablespace to hold the archive (with a retention period), creates the Flashback Data Archive, and tables are assigned to the archive (either on create or alter table statements). If different retention periods are needed, different archives must be created. Tables are assigned in whole to an archive.

### Enabling Flashback Data Archive : Example 1



## Enabling Flashback Data Archive and Usage Examples

### 1. Create the Flashback Data Archive

```
CREATE FLASHBACK ARCHIVE fla1 TABLESPACE tbs1 QUOTA 10G RETENTION 1 YEAR;
```

This has created a flashback data archive named fla1. The archive will use up to 10GB (QUOTA) from the specified tablespace (tbs1). Other tablespace(s) could be specified as needed. The archive will automatically purge historical data as it ages beyond 1 year based up the value set for RETENTION.

### 2. Grant Flashback Data Archive Privileges

```
GRANT FLASHBACK ARCHIVE ON fla1 TO scott;
```

Schema owner 'scott' now has permission to write historical data into the fla1 flashback data archive.

### 3. Begin Tracking Data Changes

```
ALTER TABLE accounts FLASHBACK ARCHIVE fla1;
```

The ALTER TABLE begins the tracking of changes to data in the accounts table. All changes to rows in the accounts table are tracked and recorded in the fla1 flashback data archive for the retention period of 1 year.

Oracle Database maintains the history in a read-only state. Users, including the schema owner and the database administrator, can query the history but cannot alter it. In addition, Oracle Database prevents any modifications to the tables that would invalidate the history.

### 4. Point-In-Time Query Example

- Use Flashback technologies to retrieve history, including:

```
SELECT ... AS OF TIMESTAMP...
```

```
SELECT ... VERSIONS BETWEEN TIMESTAMP and TIMESTAMP...
```

For the sake of this example, the data has these balance values that have changed over time:

ACCT_ID	ACCT_TIME	BALANCE
1707	01-MAY-17 13.17.00.000000 PM	1500.00
1707	02-MAY-17 09.03.00.000000 AM	25293.23
1707	02-JUN-17 15.57.00.000000 PM	134.98
1707	01-AUG-17 15.57.00.000000 PM	130900.00

Querying the data today would show only the following current balance:

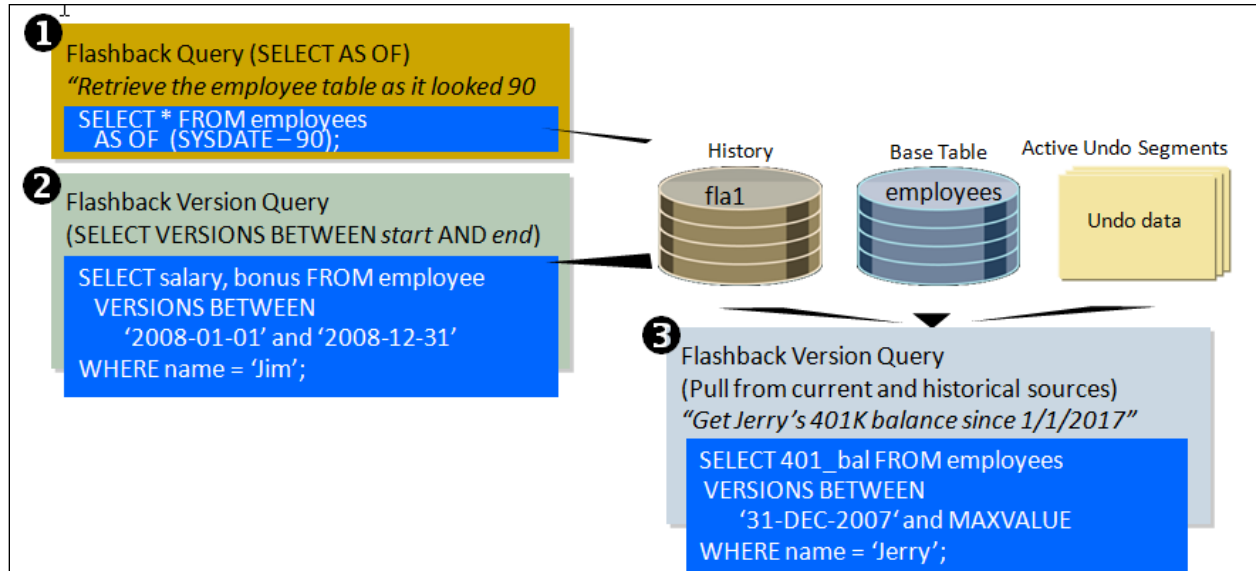
ACCT_ID	ACCT_TIME	BALANCE
1707	01-AUG-17 15.57.00.000000 PM	130900.00

However, the customer wants to know when their balance changed prior to 1 June 2017, a simple query with FDA enabled will provide this change information.

```
SELECT acct_id, acct_time, balance FROM accounts WHERE acct_time <= TIMESTAMP '2017-05-31 08:00:00' AND acct_id = 1707;
```

ACCT_ID	ACCT_TIME	BALANCE
1707	01-MAY-17 13.17.00.000000 PM	1500.00
1707	02-MAY-17 09.03.00.000000 AM	25293.23

### Additional Flashback Query Usage Examples



FDA creates history tables but does not create an index on any data column in the history table. The user can create index on any column of the history table directly. When the user issues an "as of" the optimizer may choose to use the index for the query based on the cost of the plan.

Oracle automatically partitions history tables based on a range-partitioning scheme. The compression depends whether or not the flashback archive is defined as OPTIMIZE DATA and availability of Advanced Compression. Partitioning and compression of history data is done automatically and transparently, and normally does not require administrative intervention.

## Best Practices and More Information

Below are some best practices and considerations:

- Flashback Data Archive should be enabled on a few chosen tables and not all the tables in a schema as tracking history consumes space and CPU resources, so a careful decision has to be made on the choice of the tables (and the tables tailored with the columns that need the history to be maintained)
- FDA tables should ideally be created on tables with Primary Key as a best practice so that retrieval can be based on Primary Keys and/or time/SCN ranges
- Create indexes on the History table for faster access. FDA creates history tables but does not create an index on any data column in the history table
- Try to keep the same schema for the table enabled for Flashback Archive (anticipate any future DDL changes)
- Provide enough storage so the data does not have to be exported for the retention period

Flashback Data Archive MOS Note: **FDA - Flashback Data Archive Usage and Best Practices (a.k.a. Total Recall) (Doc ID 2370465.1)**

## Conclusion

Flashback Data Archive provides many benefits over traditional application- and trigger-based methods for maintaining historic data. FDA data is tamper-proof -- no database users, not even privileged users, can change the FDA maintained data. This level of security is increasingly demanded by government regulation.

FDA lends itself to many applications. FDA allows you to configure historical data capture in a matter of minutes, provides customers with a centralized and seamlessly queryable historical data store that makes the most efficient use of all your resources – be it CPU, storage or administration time, and reduces your cost of compliance.

Use Flashback Data Archive in the same manner as other Flashback features to view or restore the data as of a time (or time range) in the past.



Oracle Corporation, World Headquarters  
500 Oracle Parkway  
Redwood Shores, CA 94065, USA

Worldwide Inquiries  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200

CONNECT WITH US

-  [blogs.oracle.com/oracle](https://blogs.oracle.com/oracle)
-  [facebook.com/oracle](https://facebook.com/oracle)
-  [twitter.com/oracle](https://twitter.com/oracle)
-  [oracle.com](https://oracle.com)

**Hardware and Software, Engineered to Work Together**

Copyright © 2019, 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.

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. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0219

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