

SecureFiles Migration

ORACLE WHITE PAPER | FEBRUARY 2019





Table of Contents

Disclaimer	1
Introduction	2
Using SecureFiles	2
Migration Techniques	3
Migration with Online Redefinition	4
Moving to SecureFiles without Migration	9
Conclusion	12

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

In Oracle Database 8i, Oracle introduced the ability to store large unstructured and semi-structured data as Large Objects (LOBs). In Oracle Database 11g Release 1, Oracle introduced a new LOB storage mechanism Oracle SecureFiles, which enables File system-like performance for LOBs. The previous LOB storage mechanism was Oracle BasicFiles.

In general, Oracle customers have typically seen performance improvements of 3x-6x, with some seeing even larger gains, by migrating to SecureFiles from BasicFiles. In order to get these gains, data stored in the older (BasicFiles) format needs to be migrated to the SecureFiles storage format.

SecureFiles represents a completely new storage paradigm for Large Objects stored in the Oracle Database. It is a complete redesign and reimplementation of LOB storage. Because of the significant changes that were required to all structures and algorithms in all levels of the database code stack to create the significant performance gains SecureFiles demonstrate, it was impossible to keep on-disk compatibility with the original LOB implementation. Thus, it is impossible to perform an “in-place” migration from BasicFiles to SecureFiles.

SecureFiles are 100% backward compatible with the various LOB APIs in the system, and migrating to SecureFiles requires no code changes to any application.

This document is intended to help Database Administrators and other technical professionals decide the best way to migrate from BasicFiles to SecureFiles.

Using SecureFiles

Requirements

To use SecureFiles, the Oracle Database parameter “compatible” must be set to 11.1.0.0.0 or higher. In addition, the tablespace the SecureFiles LOB segment (column/partition/sub-partition) resides in must be configured as “SEGMENT SPACE MANAGEMENT AUTO” (an ASSM tablespace).

Specifying SecureFiles

To specify that a LOB segment be stored as a SecureFiles, you use the “SECUREFILE” keyword as part of the LOB storage clause. The “SECUREFILE” keyword is tightly bound to the “STORE AS” phrase. Therefore, in all cases where LOB storage is being specified, “SECUREFILE” will directly follow “STORE AS” and come before any other additional keywords or specifiers that might otherwise follow “STORE AS”.

SQL DDL examples:

LOB Storage Clause

The “SECUREFILE” keyword comes after “STORE AS” and before the optional segment name:

```
LOB (c_lob) STORE AS SECUREFILE myclob
```

XMLType Storage Clause

For XMLType storage, the “SECUREFILE” keyword comes after “STORE AS” and before “CLOB” or “BINARY XML”.

CLOB storage:

```
XMLTYPE myxmltype STORE AS SECUREFILE CLOB
```

BINARY XML (BLOB) storage:

```
XMLTYPE myxmltype STORE AS SECUREFILE BINARY XML
```

VARRAY Storage

For VARRAY Storage, again, the “SECUREFILE” keyword follows “STORE AS”, and either a LOB segment name or other storage options must be specified:

```
VARRAY arr1 STORE AS SECUREFILE LOB arr1lob
```

Or alternatively:

```
VARRAY arr1 STORE AS SECUREFILE LOB (CACHE)
```

Migration Techniques

There are several ways to migrate your applications to use SecureFiles. You can migrate entire tables to use SecureFiles storage or you can migrate a partition at a time. The methods to migrate to SecureFiles are as follows:

Online Redefinition – We recommend Online Redefinition as the preferred method of migration as it allows the database and the table being migrated to be online during the migration process. There are two ways of migrating using Online Redefinition:

- a. Redefine the entire table
- b. Redefine a partition at a time

ALTER TABLE ADD PARTITION – If you do not have a need to migrate your existing data to SecureFiles, but want the performance going forward, you can create new partitions with SecureFiles storage going forward.

Creating a Partitioned Table from a Non-Partitioned Table – If your table currently is not partitioned, it is possible to leave all the current data as-is and partition your table so that new data is stored as SecureFiles.

We highly recommend one of the two Online Redefinition methods of migrating to SecureFiles.

A Simple Example

Throughout this document, we will be using a simple partitioned table with one LOB column for our examples. The example is kept simple for clarity of explaining the principles behind the various ways to migrate from BasicFiles to SecureFiles. However, the prescribed methodologies should apply to all schemas.

Sample Table

```
create table table1(  
  c_id number primary key,  
  c_lob clob  
)  
lob (c_lob) store as  
  (tablespace tbs_1)  
partition by range(c_id) (  
  partition tab1_p1 values less than (1000000),  
  partition tab1_p2 values less than (2000000),  
  partition tab1_p3 values less than (3000000),  
  partition tab1_p4 values less than (4000000));
```

Migration with Online Redefinition

The simplest and recommended method for migrating from BasicFiles (or LONGs) to SecureFiles is to use Online Redefinition.

There are many benefits to using Online Redefinition for migration:

1. The data is migrated while this system is up and available.
2. When complete all LOB data is stored in the SecureFiles storage format, and all the features and performance of SecureFiles are immediately available to your applications.
3. Addition of features such as Advanced LOB Compression, Advanced LOB Deduplication and Encryption can be done as the data is being migrated.

Online Redefinition requires free space equal to of the table/partition being migrated. However, migrating a single partition at a time can mitigate this.

Online redefinition uses materialized views to make the migration seamless. As the data is migrated, any changes to the source table will be picked up in the destination table. At the end of the migration, the source and destination table names are swapped so that all new updates and queries will operate against the new table. Upon completion, the destination table now holds the original data and can be dropped.

Migrating an entire table

This method works whether or not your table is partitioned. To migrate an entire table via online redefinition, you simply create a new table with the properties you want, in our case the LOB columns

are SecureFiles, and migrate the data to that table. On completion, you can drop this newly created table, as it will be holding the original, pre-migration, data.

During the migration more than 2x the space is required for the original table and LOBs. This is because the data will be stored in two locations, the original table and the migration table. Upon completion of the migration, the original data can be dropped and the space can be reused. If compression and deduplication are used on the destination SecureFiles columns, the amount of space required will be reduced appropriately. Please see the Administrator's Guide for more information/restrictions of online redefinition.

Let us look at a simple example of Online Redefinition of a simple table with a single LOB column.

Our source table definition:

```
create table table1(
  c_id number primary key,
  c_lob clob
)
lob (c_lob) store as
  (tablespace tbs_1)
partition by range(c_id) (
  partition tab1_p1 values less than (1000000),
  partition tab1_p2 values less than (2000000),
  partition tab1_p3 values less than (3000000),
  partition tab1_p4 values less than (4000000));
```

Our destination table definition:

```
create table table1_int(
  c_id number primary key,
  c_lob clob
)
lob(c_lob) store as SecureFile
  (tablespace tbs_sf1)
partition by range(c_id) (
  partition tab1_sf_p1 values less than (1000000),
  partition tab1_sf_p2 values less than (2000000),
  partition tab1_sf_p3 values less than (3000000),
  partition tab1_sf_p4 values less than (4000000));
```

Note: tbs_sf1 is an ASSM tablespace.

Notice two significant changes to the original table definition. They are in **bold** above.

The tablespace is in bold to note that SecureFiles must be created in an ASSM tablespace. If you are still using MSSM tablespaces for your (BasicFiles) LOB data, you will need to create or allocate ASSM tablespaces to implement the use of SecureFiles.

Performing the actual redefinition is simple. For our example, we will assume the table above, and user "Scott".

Create the destination table (table1_int from above).

1. Start the redefinition:

```
dbms_redefinition.start_redef_table(  
  'SCOTT', 'TABLE1', 'TABLE1_INT',  
  'C_ID C_ID, C_LOB C_LOB');  
-- <SCHEMA>, <Source table>, <dest_table>,  
-- <Column mapping> "source.col1 dest.col1..."
```

2. Copy the table dependents:

```
dbms_redefinition.copy_table_dependents(  
  'SCOTT', 'TABLE1', 'TABLE1_INT', 1, true,  
  true, true, false, error_count);
```

3. Finish the redefinition (swaps the table names):

```
dbms_redefinition.finish_redef_table('SCOTT', 'TABLE1',  
  'TABLE1_INT');
```

4. (Optionally) Drop the destination table:

```
DROP TABLE SCOTT.TABLE1_INT;
```

When this process completes, TABLE1's LOB column will be a SecureFiles LOB column, TABLE1_INT will have TABLE1's original definition. TABLE1_INT can now be dropped to remove the space used by the original table.

Migrating an Entire Table and Enabling Advanced Features

When performing Online Redefinition, it is possible to enable Advanced LOB Compression, Deduplication and/or Encryption (features of Oracle Advanced Compression and Oracle Advanced Security). You can do this simply by adding one or more of these advanced features to the LOB storage clause on the destination table:

```
create table table1_int(  
  c_id number primary key,  
  c_lob clob  
)  
lob(c_lob) store as SecureFile (  
  tablespace tbs_sf1  
  compress low encrypt deduplicate)  
partition by range(c_id) (  
  partition tab1_sf_p1 values less than (1000000),  
  partition tab1_sf_p2 values less than (2000000),  
  partition tab1_sf_p3 values less than (3000000),  
  partition tab1_sf_p4 values less than (4000000));
```

Note: tbs_sf1 is an ASSM tablespace.

When Online Redefinition to this destination table is complete, the CLOB column will be in stored in a compressed, deduplicated and encrypted SecureFiles column.

You can enable any advanced features, which are desired for the resulting SecureFiles data, during the migration.

All of the other migration techniques discussed in this document can also have the advanced features enabled during migration.

Online Redefinition of a Partitioned Table One Partition at a Time

If you do not have the disk capacity to migrate your entire table at one time, and the table you are migrating is partitioned, you can use online redefinition to redefine each partition one at a time.

The additional space requirement here is 1x your largest partition. Thus, the space required is significantly less than the space required to store your entire table and all LOB columns. Extra space is only required to hold one partition at a time.

In this case, we are going to operate on a partitioned table with the following definition:

```
create table scott.table1(
  c_id number primary key,
  c_lob clob
)
lob (c_lob) store as
  (tablespace tbs_1)
partition by range(c_id) (
  partition tab1_p1 values less than (1000000),
  partition tab1_p2 values less than (2000000),
  partition tab1_p3 values less than (3000000),
  partition tab1_p4 values less than (4000000));
```

The basic methodology here is to create a non-partitioned table with the SecureFiles column and to migrate each partition individually. This allows us to do the migration and only requires enough excess space to store that partition's data.

The actual steps to perform the redefinition are as follows:

```
create table scott.table1_int(
  c_id number primary key,
  c_lob clob
)
lob(c_lob) store as SecureFile
  (tablespace tbs_sf1);

declare
  i number;
begin
  for i in 1 .. 4 loop
    dbms_redefinition.start_redef_table(
      'SCOTT', 'TABLE1', 'TABLE1_INT',
```



```

        NULL, dbms_redefinition.cons_use_pk,
        '', 'tabl_p' || i);
    dbms_redefinition.finish_redef_table(
        'SCOTT', 'TABLE1',
        'TABLE1_INT', 'tabl_p' || i);
    execute immediate 'drop table scott.table1_int';
    execute immediate 'create table scott.table1_int('
        || ' c_id number primary key, '
        || ' c_lob clob) '
        || ' lob(c_lob) store as '
        || ' SecureFile '
        || ' (tablespace tbs_sf1)';
end loop;
end;
/

```

Performing the above steps will result in a table with the following definition:

```

SQL> select dbms_metadata.get_ddl('TABLE','TABLE1','SCOTT') from
dual;

CREATE TABLE "SCOTT"."TABLE1"
(
  "C_ID" NUMBER,
  "C_LOB" CLOB,
  PRIMARY KEY ("C_ID")
  USING INDEX PCTFREE 10 INITRANS 2 MAXTRANS 255
  STORAGE(INITIAL 16384 NEXT 16384 MINEXTENTS 1 MAXEXTENTS 505
  PCTINCREASE 50 FREELISTS 1 FREELIST GROUPS 1 BUFFER_POOL DEFAULT)
  TABLESPACE "SYSTEM"
  ALTER INDEX "SCOTT"."SYS_C004577" UNUSABLE ENABLE
) PCTFREE 10 PCTUSED 40 INITRANS 1 MAXTRANS 255
  STORAGE(
  BUFFER_POOL DEFAULT)
  TABLESPACE "SYSTEM"
LOB ("C_LOB") STORE AS BASICFILE (
  TABLESPACE "TBS_1" ENABLE STORAGE IN ROW CHUNK 8192 RETENTION
  NOCACHE LOGGING
  STORAGE(
  BUFFER_POOL DEFAULT))
  PARTITION BY RANGE ("C_ID")
(PARTITION "TAB1_P1" VALUES LESS THAN (1000000)
  PCTFREE 10 PCTUSED 40 INITRANS 1 MAXTRANS 255
  STORAGE(INITIAL 16384 NEXT 16384 MINEXTENTS 1 MAXEXTENTS 505
  PCTINCREASE 50 FREELISTS 1 FREELIST GROUPS 1 BUFFER_POOL DEFAULT)
  TABLESPACE "SYSTEM"
LOB ("C_LOB") STORE AS SECUREFILE (
  TABLESPACE "TBS_SF1" ENABLE STORAGE IN ROW CHUNK 8192
  NOCACHE LOGGING NOCOMPRESS KEEP_DUPLICATES
  STORAGE(INITIAL 106496 NEXT 1048576 MINEXTENTS 1 MAXEXTENTS
2147483645
  PCTINCREASE 0 BUFFER_POOL DEFAULT)) NOCOMPRESS ,
PARTITION "TAB1_P2" VALUES LESS THAN (2000000)
  PCTFREE 10 PCTUSED 40 INITRANS 1 MAXTRANS 255

```

```

STORAGE(INITIAL 16384 NEXT 16384 MINEXTENTS 1 MAXEXTENTS 505
PCTINCREASE 50 FREELISTS 1 FREELIST GROUPS 1 BUFFER_POOL DEFAULT)
TABLESPACE "SYSTEM"
LOB ("C_LOB") STORE AS SECUREFILE (
TABLESPACE "TBS_SF1" ENABLE STORAGE IN ROW CHUNK 8192
NOCACHE LOGGING NOCOMPRESS KEEP_DUPLICATES
STORAGE(INITIAL 106496 NEXT 1048576 MINEXTENTS 1 MAXEXTENTS
2147483645
PCTINCREASE 0 BUFFER_POOL DEFAULT)) NOCOMPRESS ,
PARTITION "TAB1_P3" VALUES LESS THAN (3000000)
PCTFREE 10 PCTUSED 40 INITRANS 1 MAXTRANS 255
STORAGE(INITIAL 16384 NEXT 16384 MINEXTENTS 1 MAXEXTENTS 505
PCTINCREASE 50 FREELISTS 1 FREELIST GROUPS 1 BUFFER_POOL DEFAULT)
TABLESPACE "SYSTEM"
LOB ("C_LOB") STORE AS SECUREFILE (
TABLESPACE "TBS_SF1" ENABLE STORAGE IN ROW CHUNK 8192
NOCACHE LOGGING NOCOMPRESS KEEP_DUPLICATES
STORAGE(INITIAL 106496 NEXT 1048576 MINEXTENTS 1 MAXEXTENTS
2147483645
PCTINCREASE 0 BUFFER_POOL DEFAULT)) NOCOMPRESS ,
PARTITION "TAB1_P4" VALUES LESS THAN (4000000)
PCTFREE 10 PCTUSED 40 INITRANS 1 MAXTRANS 255
STORAGE(INITIAL 16384 NEXT 16384 MINEXTENTS 1 MAXEXTENTS 505
PCTINCREASE 50 FREELISTS 1 FREELIST GROUPS 1 BUFFER_POOL DEFAULT)
TABLESPACE "SYSTEM"
LOB ("C_LOB") STORE AS SECUREFILE (
TABLESPACE "TBS_SF1" ENABLE STORAGE IN ROW CHUNK 8192
NOCACHE LOGGING NOCOMPRESS KEEP_DUPLICATES
STORAGE(INITIAL 106496 NEXT 1048576 MINEXTENTS 1 MAXEXTENTS
2147483645
PCTINCREASE 0 BUFFER_POOL DEFAULT)) NOCOMPRESS )

```

It is important to note here that even though the partition definitions changed, the *column* definition did not. This means that the default LOB storage for any future partitions added will be BASICFILE and **not** SECUREFILE, as is desired. You will have to explicitly specify SecureFiles on any columns that you might add.

As we have demonstrated, the use of online redefinition is simple and can be performed with the database online. This method also allows you to create the new partitions with advanced features such as compression, encryption or deduplication and have the migration process also perform those transformations on the data automatically. This is why we strongly recommend this method for migrating from BasicFiles to SecureFiles.

Moving to SecureFiles without Migration

Another possible method of migrating to SecureFiles is through the use of partitions. SecureFiles can be enabled on new partitions only, leaving the data in the older format as it is. So, new data will be stored in

the high performance SecureFiles format and will be able to leverage advanced features. Old data will continue to reside in the old (BasicFiles) format and can be migrated later, if desired.

ALTER TABLE ... ADD PARTITION

This method can be of particular use if your data is already partitioned by some monotonically increasing value (for example by date or via a sequence). In this case, you can simply define any new partitions you create to have SecureFiles storage for the LOB columns.

For example, consider our simple range partitioned table:

```
create table scott.table1(  
  c_id number primary key,  
  c_lob clob  
)  
lob (c_lob) store as  
  (tablespace tbs_1)  
partition by range(c_id) (  
  partition tab1_p1 values less than (1000000),  
  partition tab1_p2 values less than (2000000),  
  partition tab1_p3 values less than (3000000),  
  partition tab1_p4 values less than (4000000));
```

We can add a new partition for values less than 5000000 with a simple ALTER TABLE ... ADD PARTITION:

```
ALTER TABLE table1  
  ADD PARTITION tab1_p5 VALUES LESS THAN (5000000)  
  LOB(col3) STORE AS SecureFile (TABLESPACE tbs_sf1);
```

After this is completed, all LOB data inserted with a c_id value between 4000000 and 4999999 will be stored as a SecureFiles. This will allow you have your data going forward be stored in SecureFiles. This single largest detractor from this is the older data still has the performance characteristics of BasicFiles since the original partitions are not migrated to SecureFiles. Therefore, the performance characteristics of individual LOB accesses will be different across partitions. In addition, the table definition did not change, so you also must specify the LOB column is a SecureFiles every time you add a partition. The advantages here are you make no significant changes to your system. Assuming the data that is in the older partitions is references less and less, your application will continue to perform better as more of the data current working set is stored in SecureFiles.

Creating Partitions from a Non-Partitioned Table

If you do not have data that is already partitioned by date or some other strictly increasing or decreasing value, then you can artificially create a partitioning scheme by creating an artificial “key” and adding that key column to the existing table before the partition exchange. Here is our simple example again, but as a non-partitioned table:

```

drop table table1;
drop table table1_sf;

create table table1(
  c_id number,
  c_lob clob
)
lob (c_lob) store as
  (tablespace tbs_1);

create sequence table1_sf_key;

create table table1_sf(
  c_id number,
  c_lob clob,
  partkey$ number default 1
)
lob (c_lob) store as SecureFile
  (tablespace tbs_sf1)
partition by range (partkey$)(
  partition bf values less than (1)
    lob(c_lob) store as BasicFile,
  partition sf values less than (MAXVALUE)
);

alter table table1 add (partkey$ number default 0);

ALTER TABLE table1_sf EXCHANGE PARTITION bf WITH TABLE table1
WITHOUT VALIDATION
UPDATE GLOBAL INDEXES;

DROP TABLE table1;
RENAME table1_sf TO table1;

```

When this is complete, partition “bf” contains all the data with the LOBs stored as BasicFiles, and all data going forward will be stored as SecureFiles.

If your table is not partitioned, this could be an opportunity to add a more complicated partitioning scheme than the one provided here.

Conclusion

SecureFiles is a significant improvement in performance and features over the previous BasicFiles LOB storage mechanism. To take advantage of the performance and new features, the LOB data must be converted to the SecureFiles format.

For migration of existing data from BasicFiles to SecureFiles, we recommend online redefinition. This can be performed at the table or partition level as is necessary for your particular situation. It can be performed while the system is online and available to your applications. However, the space taken during the redefinition is roughly equivalent to 2x the storage required to store the current tables, LOBs and dependent objects.

If you are limited in the available disk space such that there is not enough space to use temporarily while the online redefinition is being performed, or if you have LOB data that is used less over time, there are other methods of migrating to SecureFiles that do not require additional space.



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



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

