
Oracle GoldenGate

Oracle GoldenGate Best Practices: Extracting from Oracle Archive log files (ALO Mode)

MOS Document ID: 1358342.1

Version 3.1

Date: November 18, 2014

Steven George
Consulting Solution Architect
Fusion Middleware Architects Team: The A-Team

Table of Contents

1	2	
2	Overview	3
3	Software Environment.....	4
3.1	Oracle GoldenGate Software install	4
4	Object ID Database Configurations	4
4.1	Object ID Database – Source Connection	4
4.2	Object ID Database – Standby Database	4
4.2.1	Standby Database as Object ID lookup only	4
4.3	Object ID Database – Object Database.....	5
5	Parameters – Archive Only	6
5.1	Parameters related to Archive Log Only Processing.....	6
5.1.1	TRANLOGOPTIONS ARCHIVEDLOGONLY	6
5.1.2	TRANLOGOPTIONS ALLOWARNEOF <seconds>	6
5.1.3	TRANLOGOPTIONS ALTARCHIVEDLOGFORMAT <string>	6
5.1.4	TRANLOGOPTIONS ALTARCHIVELOGDEST [PRIMARY] <path name>	6
5.1.5	TRANLOGOPTIONS COMPLETEARCHIVELOGONLY	6
5.1.6	TRANLOGOPTIONS COMPLETEARCHIVEDLOGTIMEOUT <seconds>	7
5.1.7	TRANLOGOPTIONS ALTARCHIVELOGDEST [PRIMARY] [INSTANCE instance_name] <path name>.....	7
5.1.8	FETCHOPTIONS, MISSINGROW REPORT	7
5.1.9	FETCHOPTIONS, NOUSESNAAPSHOT.....	7
5.1.10	FETCHOPTIONS, NOUSELATESTVERSION	8
5.1.11	STATOPTIONS REPORTFETCH	8
5.1.12	WARNLONGTRANS, USELASTREADTIME	8
5.1.13	Sample Extract Parameter file:.....	8
	Extract e_all_b	8
6	Restrictions and limitations	9

2 Overview

In systems where you don't have direct access to the source system you can now take advantage of remote data mining using Oracle Archive log files only. You can use this method of data extraction against the Source Oracle Database, Standby database, Meta data database or even BCV copy of the source database. One advantage of extracting from archive log files is that the whole process can be offloaded to another system and literally have zero impact on the source system. This method has been used to offload very high volume systems in order to load reporting databases.

This document is intended to document the best practice for setting up and configuring GoldenGate Extract to read from Oracle Archive log files only. The Object ID database can be a Source database, Standby database or even a "stub" database where the only minimal database is configured that has the object ID's from the source database

This document is organized by the different configurations for setting up archive only processing:

Oracle GoldenGate Software install - Provides an overview of GoldenGate software install.

Object ID Database, Source – The Object ID are resolved in the source database. A sqlnet net connection to the source database can be used but the Archive files that need to be processed must be local to the extract process.

Object ID Database, Oracle Standby Database – The Object ID's are resolved in a standby database that was created from the source database. Connection to the standby database can be thru a sqlnet connection, but the archive log files must be local to the extract process. If the extract resolves the object id's at startup, resolve immediate, than the standby database can be put back into managed recovery mode while extract is still running.

Object ID Database, Object database – A Object ID database can start as a standby database or a restored backup of the primary database. The only tablespaces needed in an object ID database are the system tablespace, rollback tablespace. All other tablespaces can be offline dropped before the database is opened for read/write.

Extract Parameter settings – This section covers the different parameter that need to be set for archive log processing.

Performance tuning – This section covers the recommended changes to the extract parameter file that may help extract performance.

Limitations and restrictions – This section covers restrictions and limitations for running GoldenGate in archive log only mode.

3 Software Environment

3.1 Oracle GoldenGate Software install

Oracle GoldenGate Software needs to be installing on the host per the install directions in the **Oracle® GoldenGate Oracle Installation and Setup Guide 11g Release 1 (11.1.1) E17799-01**. No special configuration needs to be completed for the reading archive only extraction.

4 Object ID Database Configurations

4.1 Object ID Database – Source Connection

This is the simplest configuration. All you need is a Oracle GoldenGate install, with the necessary Oracle install files (DLL's) and then a SQLNET connection back to the source database.

You will need some method to move the archive files to the extraction server. The files can be ftp'ed or some kind of cron job on the source can copy the files as they are created to the target.

Positives:

Any change to table structures will be picked up from the source database. Connection to the source database is only for Object ID lookup. This configuration will not be able to read current redo log member.

Negatives:

Requires a connection to the source database. Lag time, at a minimum, is the time it takes to fill a redo log file and ship it to the extract system.

4.2 Object ID Database – Standby Database

This method requires a standby database be created. In order to create the standby database you will need to read the Oracle documentation on how to create a standby database. This document assumes that the standby database has already been created. The standby database can be used in two different ways, on-line and offline. Both on-line and off-line configurations must be open in read-only mode when you start the extract process in order to resolve the object ID for the tables to extracted. If you use the resolve immediately parameters in the extract file, once the tables have been resolved you can stop the database and continue extracting from the archive log files.

Positives:

If you are resolving object id immediately, you can stop the database and continue recovery.

Negatives:

If the database has been shutdown for managed recovery and you stop or the extract process fails, you will not be able to restart the extract process until you have restarted the database in "read only" mode.

4.2.1 Standby Database as Object ID lookup only

In this configuration once you have recovered the standby database you can open the database as either “read/write” or “Read Only”. The idea behind this configuration is that you keep the database open at all times for object id lookup.

Positives:

You do not need to recover the database.

Negatives:

If tables structures are changed on the source the standby database will need to be refreshed.

4.3 Object ID Database – Object Database

The basic steps are the following:

1. Initial Database setup
 - a. Add OGG user to source DB
 - b. Add/Check database level for supplemental logging.
 - c. Add supplemental logging to tables needed in extract
2. Make backup of primary DB – core tablespaces only
 - a. This can be using RMAN, BCV splits, O.S. copy, Hot backup etc.
 - b. System tablespace
 - c. Rollback tablespace
 - d. Temp tablespace
3. Create standby controlfile
4. Copy any archive log files to target system
5. Copy Datafiles to like system
 - a. Should be same structure of you will need to add parameters to map directory names
 - b. Copy standby control file
6. Startup standby DB no mount status
 - a. Offline drop any non-system, rollback datafiles
7. Startup mount DB
 - a. Recover database (standby recovery)
8. Once recovery is until current time or last available archive file
 - a. Open database read only
 - b. If database needs more recovery before DB will open you need to check if ADV Replication is on for the source BD - Set in standby parameter file -
`replication_dependency_tracking = FALSE`
 - c. If ADV Rep is on you will need to disable it.
9. Once DB is open you can start OGG.

If the archived logs reside in a location other than the Oracle default, specify the location with the `ALTARCHIVELOGDEST` option of the `TRANLOGOPTIONS` parameter in the Extract parameter file. This parameter is supported for RAC installations in release 10 and higher.

5 Parameters – Archive Only

5.1 Parameters related to Archive Log Only Processing

Most of the parameter that relate to Archive Log Only processing are found in the parameter `TRANLOGOPTIONS` listed below.

5.1.1 `TRANLOGOPTIONS ARCHIVEDLOGONLY`

Causes Extract to read from the archived logs only, without querying or validating the logs from system views such as `v$log` and `v$archived_log`. If this parameter is specified or the database is a standby database, you will be required to position the starting position of Extract to the physical address in the log, instead of using a timestamp.

5.1.2 `TRANLOGOPTIONS ALOWARNEOF <seconds>`

Specifies the number of seconds that Extract waits for a new log file to become available before generating a warning message. Extract generates only one warning message for a given sequence number. If `ALOWARNEOF` is not specified, Extract waits for one hour by default. A value of 0 omits the warning no matter how long Extract waits.

5.1.3 `TRANLOGOPTIONS ALTARCHIVEDLOGFORMAT <string>`

(Oracle) Specifies a string that overrides the archive log format of the source database. `<string>` accepts the same specifier as Oracle's parameter `LOG_ARCHIVE_FORMAT`. Extract uses the supplied format specifier to derive the log file name

Example:

```
arch_%S.arc
```

Extract uses the supplied format specifier to derive the log file name.

5.1.4 `TRANLOGOPTIONS ALTARCHIVELOGDEST [PRIMARY] <path name>`

(Oracle) Points Extract to the archived or backup Oracle transaction logs when they reside somewhere other than the default location. Extract first checks the specified location and then checks the default location.

5.1.5 `TRANLOGOPTIONS COMPLETEARCHIVELOGONLY`

A new `COMPLETEARCHIVELOGONLY` option was added to `TRANLOGOPTIONS` to control when Extract starts processing an archived log: whether it starts immediately when the log becomes available or waits for the log to be written to disk first. This parameter applies when copying production (source) archive logs to a secondary database where they will serve as the data source. Some Oracle programs do not build the archive log from the first byte to the last byte in sequential order, but instead may copy the first 500MB, then the last 500MB, and finally the middle 1000MB, for example. If Extract begins reading at the first byte, it will abend when it reaches the break in the byte sequencing. Waiting for the whole file to be written prevents this problem.

◆ **Default in archived log only (ALO) mode:** Forces Extract to

wait for the complete archived log to be written to disk.

Override in ALO mode: Use NOCOMMPLETEARCHIVEDLOGONLY to force Extract to start processing an archived log immediately when it becomes available.

It is recommended that you leave the default value in ALO mode.

5.1.6 TRANLOGOPTIONS COMPLETEARCHIVEDLOGTIMEOUT <seconds>

(Oracle) Controls the number of seconds that Extract waits, when in COMPLETEARCHIVEDLOGONLY mode, to try again if it cannot validate that a redo log is being completely written to disk. Use this option in conjunction with the COMPLETEARCHIVEDLOGONLY option of TRANLOGOPTIONS. This option is disabled by default, and Extract will abend after ten seconds if it cannot validate that the file is being written to disk. This check is performed by reading the block header from the last block and verifying against the expected sequence number to determine if the last block has been written out. For <seconds> use any value greater than 0.

5.1.7 TRANLOGOPTIONS ALTARCHIVELOGDEST [PRIMARY] [INSTANCE instance_name] <path name>

Valid for Oracle (single instance and RAC). Points Extract to the archived logs when they reside somewhere other than the default location. Extract first checks the default location, then the location specified with this parameter. ♦ <path name> specifies the fully qualified path to the archived logs. This directory must be NFS mounted to the node where GoldenGate is running. Use that mount point for ALTARCHIVELOGDEST. INSTANCE <instance_name> applies the specified ALTARCHIVELOGDEST behavior to a specific Oracle instance. On RAC, if this option is used, you must specify the ALTARCHIVELOGDEST parameter on each node.

5.1.8 FETCHOPTIONS, MISSINGROW REPORT

This parameter provide additional information when a row is missing data in the log file. Because you are running in ALO mode, any row that is missing data will be a incomplete row and you need to know if this is happening. This parameter provides a response when GoldenGate cannot locate a row to be fetched, causing only part of the row (the changed values) to be available for processing. Typically a row cannot be located because it was deleted between the time the change record was created and when the fetch was triggered, or because the row image required was older than the undo retention specification. <action> can be one of the following: IGNORE **Ignore the condition and continue processing. This is the default.** REPORT Report the condition and contents of the row to the discard file, but continue processing the partial row. A discard file must be specified with the DISCARDFILE parameter.

5.1.9 FETCHOPTIONS, NOUSESNAPSHOT

Valid for Oracle 9i or later. The default, USESNAPSHOT, causes Extract to use the Flashback Query mechanism to fetch data needed to reconstruct operations containing LOB data, user defined data types, nested tables, and XMLType records from the undo tablespace. NOUSESNAPSHOT causes Extract to fetch the needed data from the source table. For more

information about how GoldenGate fetches data from Oracle, see the *GoldenGate for Windows and UNIX Administrator Guide*.

5.1.10 FETCHOPTIONS, NOUSELATESTVERSION

Valid for Oracle 9i or later. Use with USESNAPSHOT. The default, USELATESTVERSION, directs Extract to fetch data from the source table if it cannot fetch from the undo tablespace. NOUSELATESTVERSION directs Extract to ignore the condition if the snapshot fetch fails, and continue processing. To provide an alternate action if a snapshot fetch does not succeed, use the MISSINGROW option.

5.1.11 STATOPTIONS REPORTFETCH

REPORTFETCH returns statistics on row fetching, such as that triggered by a FETCHCOLS clause (see reference guide for details) or fetches that must be performed when not enough information is in the transaction record. NOREPORTFETCH turns off reporting of fetch statistics. The default is NOREPORTFETCH. You should not see any fetching in archive log mode.

5.1.12 WARNLONGTRANS, USELASTREADTIME

USELASTREADTIME Forces Extract to always use the time that it last read the redo log to determine whether a transaction is long-running or not. By default, Extract uses the timestamp of the last record that it read from the redo log. This applies to an Extract that is running in archive log only mode, as configured with TRANLOGOPTIONS using the ARCHIVEDLOGONLY option.

5.1.13 Sample Extract Parameter file:

Extract e_all_b

```
EXTRACT e_all_b
SETENV (ORACLE_SID=ORCL)
USERID <username> password <password>
RMTHOST <HOST_NAME>, MGRPORT 7809
RMTTRAIL ./dirdat/b1
DISCARDFILE ./dirrpt/e_all_b.dsc, APPEND
REPORTCOUNT EVERY 30 MINUTES, RATE
TRANLOGOPTIONS ARCHIVEDLOGONLY
TRANLOGOPTIONS ALTARCHIVELOGDEST PRIMARY /oracle/archive/
TRANLOGOPTIONS ARCHIVEDLOGFORMAT <format>
FETCHOPTIONS, NOUSESNAPSHOT, NOUSELATESTVERSION, MISSINGROW
REPORT
STATOPTIONS REPORTFETCH
TABLE SOURCE.*
```

6 Restrictions and limitations

The following data types can not be extracted in archive log only mode because of the requirement to fetch the missing data from the database. These data types need to be excluded from the extract process.

Data Type limitations – They vary depending on database version and GoldenGate version. Some limitations are:

- LOB records, UDT's, VARRAY, Nested Tables, XMLType, and DDL Metadata

Other restrictions:

- □Log resets (RESETLOG) cannot be done on the source database after the standby database is created.
- □To replicate DDL when Extract is in ALO mode, Extract must have permission to maintain a SQL*Net connection to the source database.
- □ALO cannot be used on a standby database if the production system is Oracle RAC and the standby database is non-RAC. In addition to both systems being Oracle RAC, the number of nodes on each system must be identical.
- ALO on RAC requires a dedicated connection to the source server. If that connection is lost, Oracle GoldenGate processing will stop.□
- Supplemental logging at the table level and the database level must be enabled for the tables from the source database.

NOTE If Extract appears to stall while operating in ALO mode, see the Oracle GoldenGate *Windows and UNIX Troubleshooting and Tuning Guide* for help with diagnosing the problem.