

3 I am keenly interested in how organizations define and manage
4 the respective roles of the Unix Systems Administrator (SA)
5 and Oracle Database Administrator (DBA). Many reasonable
6 people will maintain that it is unlikely to find one person who can
7 perform both jobs well, while others have enjoyed great success in
8 environments where there is shared responsibility for both the root
9 password and the SYSTEM password. Auditing concerns, career
10 paths, on-call schedules, and egos further cloud the issue, some-
11 times leading to the construction of brick walls where perhaps a
12 picket fence or even a chalk line on the ground would be more
13 appropriate.

14 Those who are not familiar with the latest version of Oracle's
15 flagship product, Database 10g, may be surprised to find that the
16 distinction between SA and DBA is becoming even less clear. The
17 very first document you see is called "2 Day DBA", which is tar-
18 getted to someone with no previous DBA experience, such as an
19 SA, and which promises to cover all the basic tasks one typically
20 performs in support of Oracle 10g for a small- or medium-sized
21 organization. Also, Oracle now provides many features that previ-
22 ously may have been considered the sole province of the SA, such
23 as:

- 24
- 25 • Logical disk volume manager (for all database files)
- 26 • Cluster file system and clusterware management that requires no
27 third-party software
- 28 • Lempel-Ziv compression (like gzip)
- 29 • Copy any operating system file locally or between servers
- 30 • Job scheduling capabilities, including jobs completely external to
31 the database, that vastly exceed those of cron
- 32 • POSIX-compliant regular expression handling
- 33

34 Nevertheless, one specific feature of Oracle has been in place from
35 the beginning — the archived redo log file, which is arguably the
36 single most important component of a production application for
37 the SA and DBA to jointly manage and protect. In this article, I will
38 review archived redo log files (ARLFs) from an SA perspective,
39 without the need to log in to the database or use SQL statements. I
40 will cover where archived ARLFs come from, where they go, what
41 they contain, how their use has evolved over time, and how they can
42 make or break your system. A summary checklist of issues and
43 questions will be provided that the SA and DBA can review
44 together to help ensure improved support and coverage.

45 Note that full documentation for all Oracle features mentioned
46 is freely available at the following site:

47 <http://www.oracle.com/technology/documentation/database10g.html>

49 Redo — Another Word for Journaling

51 If you understand the concept of a journaling file system, such
52 as ext3 under Linux, then you also understand most of the concept
53 behind an Oracle redo log. Every Oracle database has at least two,
54 and usually three or more, online redo log groups that exist as sep-
55 arate files in the file system (or on a raw partition), typically
56 between 10 and 100 MB in size. They are written to by the Oracle
57 logwriter process in a round-robin fashion and, in fact, are the first
58 place on disk to which any changes to the database are written,
59 even before the datafiles themselves. If the database crashes, say,
60 due to an unintended shutdown of the server, the data in these files
61 is used to help perform what Oracle refers to as "instance recovery"
62 upon the next startup, ensuring that any database transactions that

1 were committed before the crash remain intact, even if they were
2 not yet written to the permanent datafiles.

3 As updates happen in the database over time, the logwriter fills
4 up each online redo log and moves on to the next one (a log
5 switch), eventually coming back to the first one. At that point, it
6 will overwrite the previous contents, which means the amount of
7 data saved for recovery is limited by the total size of all the online
8 log groups. If there is a crash involving the physical loss of a
9 datafile due to disk failure or accidental erasure, for example, then
10 the complete sequence of redo log contents, going back to the last
11 backup of the datafile, will be needed to perform what Oracle refers
12 to as “media recovery” and ensure that no data was lost. To accom-
13 plish this, most databases run in ARCHIVELOG mode, which
14 causes the logwriter not to overwrite any online redo log until its
15 contents have been archived (copied) to some other disk-based
16 location.

17 In the beginning, this was the sole purpose of ARLFs — to pro-
18 vide for restoration of a database backup up to the time of the most
19 recent ARLF, in the event of disk failure or even a complete site-
20 wide failure. The main concern of the SA was two-fold:

- 21
- 22 1) To make sure that all ARLFs were reliably backed up together
23 with the core database files, usually to tape or another off-server
24 location.
- 25 2) To monitor the directory where ARLFs were stored so it didn’t
26 run out of space.
- 27

28 If, during normal operation, the filesystem where the ARLFs were
29 stored fills up, the logwriter cannot overwrite the oldest unarchived
30 online redo log file, bringing the entire database to a screeching
31 halt until the space problem was addressed.

32 Here is a sample file listing (abbreviated `ls -lt` output) show-
33 ing an example of the kind of behavior that ARLFs can exhibit.
34 Note that in this two-hour window, the rate of disk consumption has
35 suddenly gone from roughly half a megabyte per hour to nearly 100
36 MB per minute — a 12,000-fold increase!

```
37  
38 20968960 Jan 30 18:10 binc02_0001_0000000097.dbf  
39 20970496 Jan 30 18:10 binc02_0001_0000000096.dbf  
40 20967936 Jan 30 18:10 binc02_0001_0000000095.dbf  
41 20969984 Jan 30 18:10 binc02_0001_0000000094.dbf  
42 20971008 Jan 30 18:10 binc02_0001_0000000093.dbf  
43 20970496 Jan 30 18:09 binc02_0001_0000000092.dbf  
44 20970496 Jan 30 18:09 binc02_0001_0000000091.dbf  
45 20971008 Jan 30 18:09 binc02_0001_0000000090.dbf  
46 20971008 Jan 30 18:09 binc02_0001_0000000089.dbf  
47 628736 Jan 30 17:10 binc02_0001_0000000088.dbf  
48 330240 Jan 30 16:10 binc02_0001_0000000087.dbf
```

49
50 What’s more, this could be either normal and expected, or a symp-
51 tom of something gone terribly wrong, such as a software bug or
52 someone making a data alteration they should not be making, at
53 least not at this time of day.

54 **Finding the Way to Oracle Home**

55 To stay on top of ARLF generation, you need to know where
56 they are. Let’s suppose that you have taken over a running database
57 system and that there is no one to ask about how it was configured.
58 The first thing you need to find is the Oracle “home”, the unique
59 base directory used for each installation of Oracle software on the
60 server. To support concurrent installation of different versions and
61 different products, there can be many Oracle homes on the same
62

1 box. The best place to look is in the oratab file, which is created as
2 part of the Oracle software installation process. The file will be
3 located as follows:

```
4  
5 Solaris: /var/opt/oracle/oratab  
6 AIX, HP, Linux and Tru64: /etc/oratab
```

7
8 The entries in this plain-text file consist of three colon-separated
9 fields as follows:

```
10  
11 ORACLE_SID:ORACLE_HOME:{Y|N }
```

12
13 For each database instance configured on the server, the entry in
14 oratab shows its name, Oracle home location, and whether it should
15 be started and stopped by the Oracle-supplied dbstart and dbshut
16 scripts whenever they are run. Another, less reliable way to dis-
17 cover the candidate locations for ORACLE_HOME is to examine
18 the shell startup scripts (.profile, .bashrc, etc.) for the “oracle” user
19 account, which typically will have at least one setting for an ORA-
20 CLE_HOME environment variable.

21 Once you have a list of ORACLE_HOME locations for the box,
22 set your \$ORACLE_HOME shell environment variable to each
23 value in turn and look in each location for a subdirectory named
24 dbs. In this subdirectory, along with a few other default files, you
25 will typically find files of the form initSID.ora, where SID is the
26 specific database instance name corresponding to the first field in
27 the oratab file.

28 For example, if the oratab entry is binc02:/u01/app/oracle/prod-
29 uct/9.2.0:Y, then the SID is “binc02”, and the file you are looking
30 for will be named “initbinc02.ora”. This parameter file (PFILE) is
31 essential for starting up any instance of Oracle. Of particular rele-
32 vance is the parameter that indicates where ARLFs are to be gener-
33 ated. In the simplest case, you will be able to examine the PFILE
34 for each instance and find one or more entries that look like this:

```
35  
36 log_archive_dest_1 = \  
37 'LOCATION=/u01/app/oracle/admin/binc02/arch MANDATORY'
```

38
39 Up to 10 log archive destinations can be defined. If you find no cur-
40 rent files in the directory specified by LOCATION=, then probably
41 the database instance has not been started in a long time, or it is in
42 NOARCHIVELOG mode.

43 Let’s examine three common exceptions to this simple process
44 for finding the ARLF location, especially those due to newer ver-
45 sions of Oracle. In addition to a standard plain-text PFILE, starting
46 with version 9i, a more flexible binary version (known as a server
47 parameter file, or SPFILE) can be used instead of a PFILE. Or, the
48 PFILE can have a single entry pointing to a SPFILE, such as:

```
49  
50 SPFILE='/opt/oracle/oradata/bin1r/spfilebin1r.ora'
```

51
52 If the dbs directory contains a PFILE pointing to an SPFILE as
53 shown here, or if it contains an SPFILE directly, with a name
54 matching spfileSID.ora or spfile.ora, then you need to examine the
55 SPFILE for the log_archive_dest_n parameter. Fortunately,
56 although this file is binary, you can easily examine its contents
57 without logging in to the database by using the strings command,
58 as this example using a SID of bin1r shows:

```
59  
60 % strings spfilebin1r.ora | grep log_archive_dest_1  
61 *.log_archive_dest_1='LOCATION=/opt/oracle/admin/bin1r/arch'
```

62

1 A second exception has to do with the dbs directory itself. If you
2 find an ORACLE_HOME without one, then it represents an instal-
3 lation of a product that doesn't support a database, such as a client-
4 software-only installation, or a cluster-ready services (CRS)
5 installation. A third exception shows up when you discover an
6 instance name or ARLF destination that begins with a plus sign "+"
7 as in these two examples:

```
8  
9 oratab: +ASM:/u01/app/oracle/product/10.1.0:Y  
10  
11 initbinc03.ora: SPFILE=+BINC_LAKE/BINC03/PARAMETERFILE/spfile.262.1
```

12
13 In this case, we are encountering a new feature available under
14 Oracle 10g, Automated Storage Management (ASM). This is a disk
15 volume manager that provides striping and mirroring for essentially
16 all Oracle database files, using disk partitions, third-party volumes,
17 or space on certified NAS devices. Unfortunately, examining and
18 manipulating the contents of an ASM volume requires logging into
19 the database, so it will remain outside the scope of this article.

20 There are two other much less common exceptions to note as
21 well — non-standard (S)PFILE locations used when starting the
22 database, and parameter changes made in the running database that
23 aren't reflected in the (S)PFILE.

24 Redo Log File Transport Using Oracle Net

25
26 Beginning with Oracle version 8, a new type of database known
27 as a standby database was provided. A physical standby database is
28 basically a "live" backup, usually kept current with transactions no
29 more than a few hours behind the primary database and ready to go
30 live in a recovery situation in a matter of minutes. In Oracle 9i, a
31 logical standby type was added, and both types are now known by
32 the name Data Guard. The two hallmarks of a typical Data Guard
33 database are that it is on a physically separate server from the pri-
34 mary, and transactions are copied into the standby from the primary
35 database solely by way of reading the contents of the ARLFs gen-
36 erated on the primary (which has the benefit of not adding any pro-
37 cessing load to the primary database, as other types of data
38 replication typically would).

39 Although ARLFs can be copied and applied at the standby
40 server under the control of Unix-level scripts, use of Oracle's auto-
41 matic log transport mechanism is recommended. The SA can deter-
42 mine whether and how this is configured by looking in the
43 (S)PFILE for a log_archive_dest parameter that references a
44 SERVICE instead of a LOCATION, for example:

```
45  
46 log_archive_dest_3 = 'SERVICE=binc02dg'
```

47
48 The service is an Oracle Net service, and the details can usually be
49 determined by using the Oracle tnsping utility (found in \$ORA-
50 CLE_HOME/bin directory):

```
51  
52 % tnsping binc02dg | grep DESCRIPTION  
53 Attempting to contact (DESCRIPTION = (ADDRESS_LIST = \  
54 (ADDRESS = (PROTOCOL = TCP)(HOST = tree)(PORT = 1521))) \  
55 (CONNECT_DATA = (SERVICE_NAME = binc02dg)))
```

56
57 The output shows that the remote destination for ARLFs defined by
58 this parameter is at a network address named "tree", where the
59 Oracle Net Listener is using the standard network port 1521. On the
60 remote host "tree", you can look (using the techniques described
61 above) for an ORACLE_HOME and a SID of "binc02dg"; in the
62 (S)PFILE for this instance, you should find the following setting:

```
1
2 standby_archive_dest = '/u01/app/oracle/admin/binc02dg/stbyarch'
```

3
4 This indicates that the standby database “binc02dg” on host “tree”
5 will store the copies of ARLFs it receives from the primary to the
6 directory shown. From there, the Data Guard update process will
7 apply the transactions to the standby database. Note that if a logical
8 standby is in ARCHIVELOG mode, it will in turn generate its own
9 local ARLFs — a good reason to have the parameters
10 log_archive_dest_n and standby_archive_dest set to different
11 directories in the standby instance. Also note that there is not a one-
12 to-one correspondence between the two sets of ARLFs — it is not
13 unusual for the local ones to consume much more space than the
14 received copies.

15 To avoid data loss due to the primary database being a single
16 point of failure, Oracle provides a Data Guard configuration where
17 transactions won't complete at the primary database until they are
18 first written to the standby ARLF. Obviously, you will want an
19 extremely reliable network and server configuration before putting
20 this restriction on your primary database! In a more typical sce-
21 nario, you will accept slightly increased risk by simply making sure
22 that ARLFs are generated frequently, say every 15 or 30 minutes,
23 which will become the maximum possible time period of lost data
24 in the event of an unrecoverable failure of online redo logs at the
25 primary site. ARLFs are normally generated only when the online
26 redo log fills up and a log switch must occur; however, the follow-
27 ing parameter can be used on the primary database to automatically
28 ensure that a log switch happens at set intervals whether needed or
29 not:

```
30
31 archive_lag_target = 900 # value is in seconds
```

32 Security, Retention, and Performance

33 Issues for Archived Redo Log Files

34
35 When disaster recovery was the only purpose of ARLFs, you
36 could delete and forget all that were older than the last successful
37 database backup. In current versions of Oracle's product, it is now
38 possible to use Log Miner to extract every single database update
39 directly from the ARLFs, without any continued access to the origi-
40 nal database or its passwords. Even transactions that did not com-
41 plete, or were rolled back, at the primary are available by mining
42 the ARLFs. For this reason, all ARLFs and copies should be main-
43 tained as securely as the primary database. Because they are rela-
44 tively small and typically copied to several locations, it is easy to
45 lose track of them. If appropriate, your organization might consider
46 encrypting them both in transport and in storage for added protec-
47 tion.

48 The ability to mine the contents of ARLFs also makes them use-
49 ful for auditing or troubleshooting activities, and you may decide to
50 keep them around for a year or more for this reason. For example,
51 one time I noticed a pattern over several weeks of a huge surge in
52 ARLF generation at a certain time early in the morning. Yet there
53 was no immediate evidence of any users performing any unusual
54 transactions. By mining the available ARLFs for the relevant time
55 period, I discovered the ultimate problem — an application bug,
56 coupled with Oracle's sophisticated handling of transaction isola-
57 tion levels, was causing an attempted update to a table to be par-
58 tially rolled back and re-attempted many hundreds of times.
59 Without the ARLFs, it would have been nearly impossible to iden-
60 tify the exact user-initiated steps and database tables that led to
61 uncovering the bug.

62

1 The “huge surge” in ARLF generation just described should
2 always prompt further investigation by the SA and DBA. Even if
3 the cause is known and expected, be aware that rapid online redo
4 log switches and generation of ARLFs consumes a large amount of
5 system resource and can noticeably impact response time for other
6 users of the system.

7 A classic example that all SAs for Oracle systems should under-
8 stand goes like this — suppose the database appears to be “hung”
9 because of some kind of runaway transaction that is generating
10 ARLFs at a high rate and is under pressure from users; the decision
11 is made to bounce it immediately. The SA may end up killing the
12 Oracle processes at the Unix level or even rebooting the server to
13 ensure a clean restart. This is perfectly acceptable to the Oracle
14 database, and no committed data will be lost due to automatic
15 instance recovery. When the database starts up again, the “hung”
16 transaction, which may have been, say, two-thirds complete, must
17 be rolled back to guarantee database integrity. However, each data-
18 base block that is rolled back is in turn logged, resulting in another
19 huge surge of ARLF generation. This activity can actually be sev-
20 eral times longer and consume several times more disk space than
21 the original interrupted transaction, so in this scenario the users are
22 actually out of service longer than if nothing had been done.

23 **Oracle’s Newer Automated Management** 24 **Features**

25 Recognizing that many aspects of managing ARLFs have his-
26 torically been addressed at the Unix level — external to the data-
27 base — Oracle has implemented several new features to relieve the
28 SA of this burden. Even if you are not yet ready to migrate to
29 Oracle 10g, knowing what is available can help you be better pre-
30 pared when the time comes.

31 Traditionally, the names of the ARLFs are generated according
32 to a format specification in the (S)PFILE, such as the following:

```
33 log_archive_format = binc02_%T_%S.arc  
34
```

35 where %T is the instance thread number, and %S is the log
36 sequence number. This can cause a conflict when a new incarnation
37 of the database is created, which resets the sequence number back
38 to 1. Using the wrong ARLF, because it has the same name as an
39 older one from a previous incarnation, can be disastrous during a
40 recovery effort. In version 10g, an additional format variable (%r)
41 is now required to be part of the name to ensure unique file names
42 across multiple incarnations of the database.

43 Also in 10g, the Flash Recovery Area (FRA) feature can be used
44 to automate most disk-based backup and recovery functions.
45 Within specified space and retention limits, Oracle will automati-
46 cally manage the placement and eventual deletion of ARLFs. The
47 FRA is also integrated with the Oracle Managed Files (OMF) fea-
48 ture, which will automatically generate unique names for most
49 types of database files, including ARLFs. If the following parame-
50 ter is set in the (S)PFILE, then ARLFs will be generated automati-
51 cally in the FRA location shown:

```
52 db_recovery_file_dest='/opt/oracle/oradata/flash'  
53
```

54 and the generated ARLF names will typically look like this:

```
55 o1_mf_1_104_0zx14krz_.arc  
56  
57  
58
```

1 Since files in the FRA are subject to automatic deletion, you may
2 want to ensure an extra copy is maintained outside of the FRA,
3 using parameters as shown below, where destination 1 is outside
4 the FRA and destination 2 is in the FRA:

```
5  
6 *.log_archive_dest_1='LOCATION=/opt/oracle/admin/bin1r/arch'  
7 *.log_archive_dest_2='LOCATION=USE_DB_RECOVERY_FILE_DEST'
```

8
9 The Oracle backup and recovery tool, RMAN, is also integrated
10 with the FRA and can take full advantage of automated manage-
11 ment of ARLFs when performing its tasks.

12 Clusters and the Other Kind of Log File

13
14 Two additional topics related to ARLFs merit a brief mention.
15 In a Real Application Cluster (RAC) configuration, two or more
16 Oracle instances running on different servers share a common data-
17 base. Each instance generates its own sequence of ARLFs, so it is
18 important to identify all the destinations involved from the SPFILE
19 for the cluster. Typically, the ARLFs will be copied to at least one
20 common, shared location, using ASM or a cluster file system. Each
21 instance uses a different thread number, so the source of each file
22 will be normally distinguishable by the thread number embedded in
23 its name.

24 When all else fails, good old-fashioned Oracle runtime log files
25 (the human-readable ones) are an invaluable source of information
26 for both normal operation and troubleshooting. To find the primary
27 log file for an instance, look for the entry in the (S)PFILE similar to
28 this:

```
29  
30 background_dump_dest = /u01/app/oracle/admin/binc02/bdump
```

31
32 Among the other trace files in this location, a special trace file
33 called the “alert” log, containing all the core messages for the
34 instance, will be found with a name of alert_SID.log. Among other
35 things, this file will show you all non-default parameters from the
36 (S)PFILE that was used to start the instance. It also logs the cre-
37 ation time and success of each ARLF. For example, if Oracle is
38 shipping ARLFs to a standby database that happens to be out of
39 service at the moment, you will see error messages such as the fol-
40 lowing in the primary database alert log:

```
41  
42 Wed Feb 2 16:49:33 2005  
43 Errors in file /u01/app/oracle/admin/binc02/bdump/binc02_arc1_1726.trc:  
44 ORA-12514: TNS:listener could not resolve SERVICE_NAME given in  
45 connect descriptor
```

46
47 By examining the referenced trace file, binc02_arc1_1726.trc, you
48 find the following:

```
49  
50 Error 12514 attaching to destination LOG_ARCHIVE_DEST_3 standby  
51 host 'binc02dg'  
52 Heartbeat failed to connect to standby 'binc02dg'. Error is 12514.
```

53
54 The messages found in the alert log and other trace files will be
55 your primary source of information when working with other sup-
56 port personnel.

57 Conclusion

58
59 Whether you work closely with — or worlds apart from — your
60 DBA counterpart, it behooves the SA supporting an Oracle data-
61 base server to understand the impacts of the database on the rest of
62 the system and know where to find key information about how

- 1 things are configured, especially under the newest version of
2 Oracle. You may find the following brief checklist handy as a start-
3 ing point for further discussion with the DBA in your shop:
4
- 5 1. Are oratab files up-to-date and in the standard location for all
6 instances?
 - 7 2. Are all parameter files (PFILE/SPFILE) in the standard location,
8 backed up, and contents documented?
 - 9 3. Are there any user-managed scripts at the Unix level for copying
10 ARLFs to remote locations?
 - 11 4. If using Flash Recovery, are additional copies of ARLFs stored
12 outside the Flash Recovery Area?
 - 13 5. If using ASM, are additional copies of ARLFs generated outside
14 of ASM storage?
 - 15 6. Do all ARLFs need to be kept secure (encrypted)?
 - 16 7. How long do ARLFs need to be saved?
 - 17 8. Are ARLFs generated frequently enough (minimum every 15-30
18 minutes), or do you implement maximum protection mode for
19 Data Guard standby databases?
 - 20 9. Are log file names standardized and forward compatible with
21 10g?
 - 22 10. Are logical standby local ARLFs stored in a different directory
23 from the ARLFs received from the primary database, or alter-
24 natively, can the logical standby run in NOARCHIVELOG
25 mode?
 - 26 11. Are database maintenance and update activities involving large
27 transactions fully tested with respect to the quantity of ARLF
28 generated before going into production?
 - 29 12. Are Oracle alert log files regularly monitored for error mes-
30 sages?

**Author — Please provide 2-3 sentence bio
for publication. Thank you!**