



Preventing, Detecting and Repairing Block Corruption in 12c

Contents

Introduction	3
Types of Corruption	3
Physical Corruption	3
Logical Corruption	3
Intra block Corruption	4
Inter block Corruption	4
How Corruption Occurs at First Place	4
Risks	4
MAA: Data Protection for all Databases	5
Tools	5
DB Verify	5
ANALYZE.. VALIDATE STRUCTURE	5
DB_BLOCK_CHECKING	6
RMAN (BACKUP VALIDATE, RESTORE VALIDATE, VALIDATE)	6
Block Media Recovery	6
Data File Media Recovery	6
Block Corruption Parameter	6
High Availability Features in Oracle Database 12c	6
Application Continuity	7
ASM Disk Scrubbing	7
Automatic Block Repair– High Availability	7

Introduction

When data is not readable by oracle or it is not in a recognized Oracle Database format, or its contents are inconsistent, it is corrupted. Data block corruption can damage internal Oracle control information or application and user data, leading to loss of services or partial/full outages

Block corruptions may affect only a single block or a large portion of the database essentially. Oracle provides a complete set of technologies to prevent and mitigate block corruptions.

Corruption can be due to hardware or software issues, and prevention is better way to address the issue than cure. While we cannot prevent all block corruptions, there is comprehensive set of data protection solutions that we can implement to address most of them:

Oracle Data Guard and Active Data Guard

Data Recovery Advisor

Oracle Flashback

Oracle Recovery Manager (RMAN)

Automatic Diagnostic Recovery (ADR)

Oracle Secure Backup

The MAA Advisor component of Oracle Enterprise Manager Grid Control

Ex a data Storage

Types of Corruption

Physical and Logical Corruptions

Data corruption can manifest itself as a physical or a logical corruption:

Physical Corruption: This is when block has an invalid check sum or header, or when the block contains all zeroes. When that occurs, the database will call the block as a corrupt block, regardless of its content. A physical corruption is also called a media corruption.

Logical Corruption: When a data block has a valid check sum, etc., but the block contents are logically inconsistent, it is

logically corrupt. Logical block corruption can also occur when the structure below the block header is corrupt. In this case, the block check sum is correct but the block structures may be corrupt.

Intra block Corruption: The physical or logical corruption that occurs in the block itself.

Inter block Corruption: The corruption that occurs between blocks and is a logical one.

How Corruption Occurs at First Place

The sequence for I/O goes like:



Corruption maybe an usual occurring but shouldn't be an recurring problem. The best way to avoid it is to prevent it.

Hardware failures or bugs in any layer can result in corrupt data being written to disk, however it is reported as written to oracle.

Risks

Data can be corrupted anywhere and anytime...and can be undetected unless touched

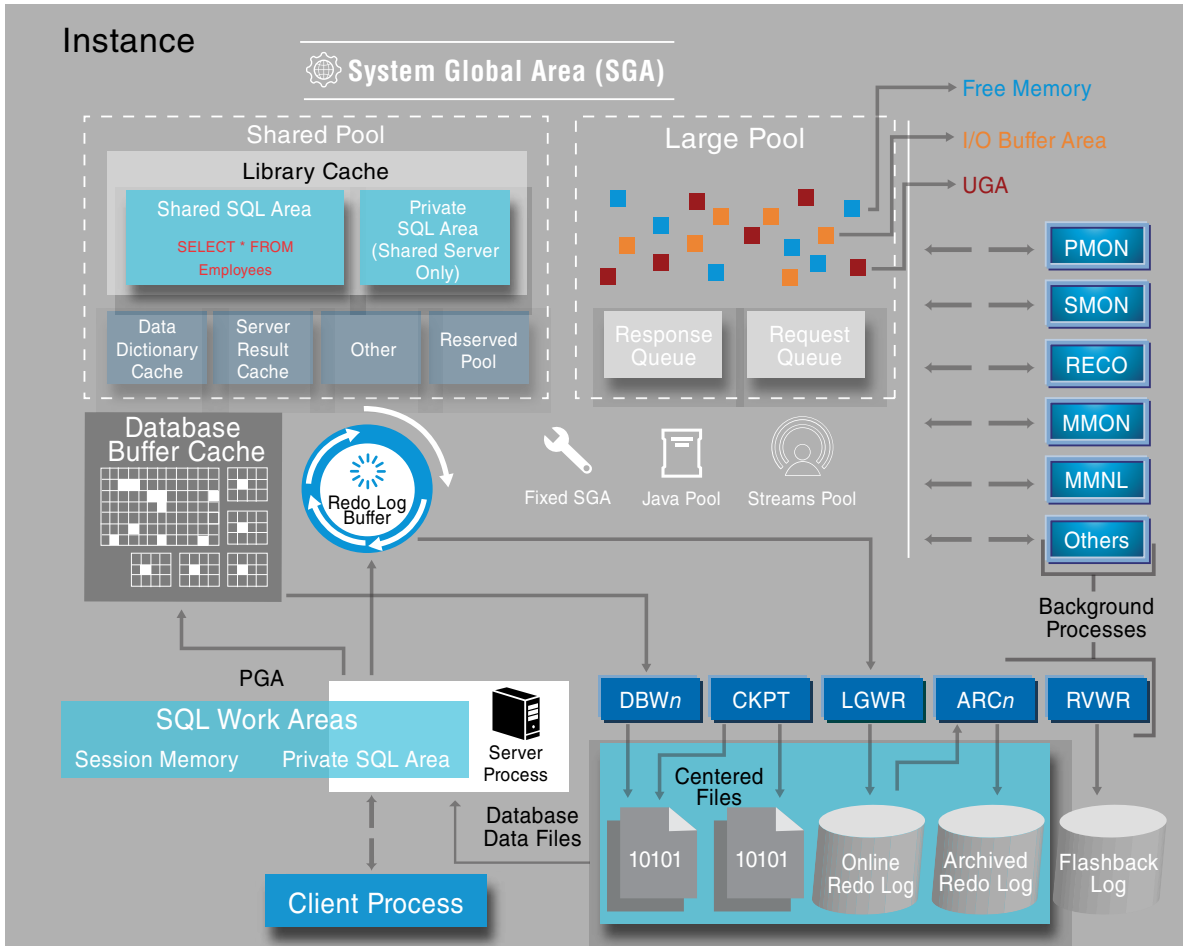


Fig. No.1

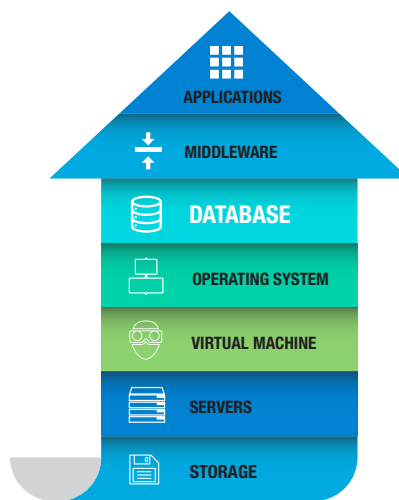


Fig. No.2

Checksum is not sufficient




Fig. No.3

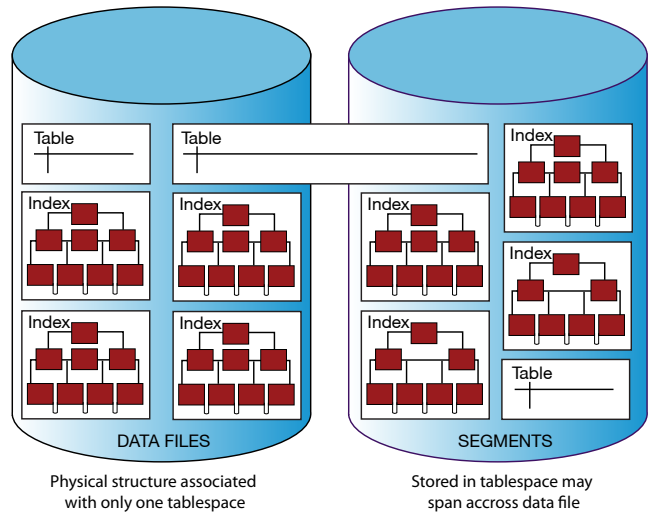


Fig. No.4

Backups and DR without valid ationisen or mous risk .Do not guarantee working or meeting recovery SLAs

MAA: Data Protection for all Databases

Oracle MAA prescribes four reference architectures that provide progressively higher levels of availability and data protection.

Each architecture is based upon a common infrastructure optimized for Oracle Database.



Fig. No.5

The MAA recommendation to achieve the most comprehensive data corruption prevention and detection is to use

1. Oracle Data Guard with physical standby databases and configure the DB_BLOCK_CHECKING, DB_BLOCK_CHECKSUM, and DB_LOST_WRITE_PROTECT initialization parameters on the Data Guard primary and standby databases.

2. Oracle Automatic Storage Management (Oracle ASM) that provides disk mirroring to protect against disk failures. Once the corruption is detected, Oracle Data Guard, block media recovery, and data file media recovery can be used to recover the data.
3. Oracle Flashback Technologies to undo the Database-wide logical corruptions caused by human or application errors.

Tools

Tools are also available for proactive validation of logical data structures.

DB Verify

DB Verify is an external utility that allows validation of offline and online data files. In addition to offline data files it can be used to check the validity of backup data files.

```
tst db>dbvfile=C:\Oracle\oradata\test\system dbf
feedback=10000blocksize=8192
```

ANALYZE .. Validate Structure

The ANALYZE command can be used to verify each data block in the analyzed object. If any corruption is detected, rows are added to the INVALID_ROWS table. sql>ANALYZE TABLE scott.emp VALIDATE STRUCTURE

DB_BLOCK_CHECKING

When the DB_BLOCK_CHECKING parameter is set to [TRUE|HIGH] Oracle performs a data check in the block for self-consistency. Unfortunately block checking can add between 1 and 10% overhead to the server.

Oracle recommend setting this parameter to [TRUE|HIGH] if the overhead is acceptable. Allowable values include [OFF|FALSE], LOW, MEDIUM, [HIGH|TRUE]

RMAN (BACKUP VALIDATE, RESTORE VALIDATE, VALIDATE)

```
RMAN>BACKUP VALIDATE DATABASE ARCHIVE LOG ALL;
```

```
RMAN>BACKUP VALIDATE CHECK LOGICAL DATABASE ARCHIVE LOG ALL;
```

```
RMAN>RESTORE DATABASE VALIDATE;
```

```
RMAN>RESTORE ARCHIVE LOG ALL VALIDATE;
```

V\$ DATABASE_BLOCK_CORRUPTION view displays blocks marked corrupt by database components such as RMAN, ANALYZE, and SQL queries

Block Media Recovery

We can use block media recovery to recover one or more corrupt data blocks within a data file. It provides the following advantages over data file media recovery:

- Lowers the meantime to recover (MTTR) because only blocks needing recovery are restored and recovered
- Enables affected data files to remain online during recovery

Block media recovery is most useful for physical corruption problems that involve a small, known number of blocks. Block-level data loss usually results from intermittent, random I/O errors that do not cause widespread data loss, and memory corruptions that are written to disk.

Block recovery can be performed only on blocks that are marked or detected as corrupt by corruption check.

Database automatically attempts to perform block media recovery if the database is associated with a real-time query physical standby database. Primary database uses good copies of block on standby to repair the corrupt block.

Manual recovery can be done using RECOVER... BLOCK. RMAN perform search in order of -

- Real-time query physical standby database
- Flashback logs and then blocks in full
- Level 0 incremental backups

High Availability Features in Oracle Database 12c

Application Continuity

Oracle Database 12c provides the improved user-experience for applications using Application Continuity.

Automatic block repair is performed if the following parameters are set on standby -

- The LOG_ARCHIVE_CONFIG parameter is configured with a DG_CONFIG list and a LOG_ARCHIVE_DEST_n parameter is configured for the primary database with the DB_UNIQUE_NAME attribute
or
- The FAL_SERVER parameter is configured and its value contains an Oracle Net service name for the the Primary database

RMAN command RECOVER CORRUPTION LIST recovers all blocks detected corrupt from the V\$ DATABASE_BLOCK_CORRUPTION

Data File Media Recovery

Data file media recovery repairs a lost or damaged current data file or control file. It can also recover changes lost when a table space goes offline without the OFFLINE NORMAL option.

Block Corruption Parameter

Configure the following on the Data Guard primary database:

- DB_BLOCK_CHECKSUM=FULL
- DB_BLOCK_CHECKING=FULL or MEDIUM
- DB_LOST_WRITE_PROTECT=TYPICAL
- Enable Flashback Technologies for fast point-in-time recovery from logical corruptions most often caused by human error and for faster in statement of the news standby database following fail over.

Configure the following on the Data Guard standby database:

- DB_BLOCK_CHECKSUM=FULL
- DB_BLOCK_CHECKING=FULL or MEDIUM
- DB_LOST_WRITE_PROTECT=TYPICAL
- Enable Flashback Technologies for fast point-in-time recovery from logical corruptions most often caused by human error.
- Use Active Data Guard to enable Automatic Block Repair (Active Data Guard release 11.2 and later).

Understanding the different types of corruption helps you recover the failure easily.

- Replays in-flight work on recoverable errors
- Masks many hardware, software, network, storage errors and outages
- Improves end user-experience and developer productivity

When replay is successful, Application Continuity masks many recoverable database outages from the applications and the users. It achieves the masking by restoring the database session, the full session (including session states, cursors, variables), and the last in-flight transaction (if there is one).

If the database session becomes unavailable due to a recoverable error, Application Continuity attempts to rebuild the session and restore any open transactions to the correct states.

If the transaction is successfully committed and does not need to be re-executed, the successful status is returned to the application.

If the replay is successful, the request continues safely, with no risk of duplication.

If the replay is not successful, the database rejects the replay and the application receives the original error. To be successful, the replay must return to the client the exact same data that the client received previously in the request, which the application potentially made a decision on.

ASM Disk Scrubbing

When using ASM mirroring (normal or high redundancy), there are more than one copy of an allocation unit (AU). The AU and its copies should be in sync. When logical corruption creeps in, there is higher response time from the disk affected by the corruption. In Oracle Database 12c ASM we can use a new command called SCRUB to weed out the logical corruption. This command repairs the logical corruption by reading the data from the mirror copies. Here is how we can use to repair the disk group DATA:

```
SQL>alter disk group data scrub repair; Disk group altered.
```

Again, as with the previously described operations involving large movements of data between disks, we can control how much resource this operation will take by using a special clause called "power".

However, instead of a number, this parameter expects values from the list: LOW, HIGH, MAX and AUTO. A power of MAX will consume most resources to complete the operation faster, but may affect other operations in the system. Here is how:

```
SQL>alter disk group data scrub power max;
```

The power value of AUTO lets ASM choose the best power value depending on the system resources available. This is also the default option. If the load on the system I/O is very high, the scrubbing operation is not performed by Oracle since it will just make the I/O response even worse. To force the scrubbing even under those circumstances FORCE clause is used:

```
SQL>alter disk group data scrub repair power max force;
```

But scrubbing is not just for the entire disk group; we may choose to scrub a single disk as well. This is helpful if we want

to break up the activities to one disk at a time. Here is how we do it for a disk:

```
SQL>alter disk group data scrub disk data_0000 repair power max force; Disk group altered.
```

We can even repair a specific file. This is particularly useful when we want to make sure important files such as system or sysaux data files and vital application related files are scrubbed first.

```
SQL>alter disk group data scrub file' +DATA/CONA/DATAFILE/USERS.271.824035767' repair; Disk group altered.
```

If we want to merely check the presence of the logical corruptions, not actually fix them, we need to just omit the keyword "repair" in the command.

```
SQL>alter disk group data scrub file' +DATA/CONA/DATAFILE/USERS.271.824035767';
```

This will report the logical corruptions and not fix them.

Automatic Block Repair – High Availability

Block-level data loss usually results from intermittent random I/O errors, as well as memory corruptions that get written to disk. When Oracle Database reads a block and detects corruption, it marks the block as corrupt and reports the error to the application. No subsequent read of the block will be successful until the block is recovered manually, unless you are using Active Data Guard.

Active Data Guard automatically performs block media recovery that is transparent to the application. Active Data Guard repairs physical corruption on a primary database using a good version of the block retrieved from the standby. Conversely, corrupt blocks detected on the standby database are automatically repaired using the good version from the primary database.

Physical corruption on an active standby database is also detected and automatically repaired even in cases where a block has never been changed at the primary database or read by applications running at the standby. This is done by enabling Data Guard lost-write protection at both primary and standby databases; a standard best practice for detecting silent corruption resulting from transactions that use stale data. Lost-write protection has a secondary benefit of dramatically increasing the overall level of validation for physical corruption performed at a standby database. Lost-write validation occurs at the standby database for every block that is read at the primary, whether or not the data is changed. Reading the standby version of the block in this manner triggers additional checks for physical block corruption to detect faults that occur only at the standby database and not at the primary.



Shilpa Rajpurkar

DBA Lead Consultant - Oracle DBA

Shilpa Rajpurkar is an Oracle DBA. Currently working as a DBA Lead Consultant, she has been working with Mphasis since October 2009. With a total experience of 14 years in IT industry, she has worked in Insurance, Retail and telecom domain.

Currently working on 12c, Shilpa has worked on most of the versions of oracle.

About Mphasis

Mphasis (BSE: 526299; NSE: MPHASIS) applies next-generation technology to help enterprises transform businesses globally. Customer centricity is foundational to Mphasis and is reflected in the Mphasis' Front2Back™ Transformation approach. Front2Back™ uses the exponential power of cloud and cognitive to provide hyper-personalized ($C = X2C_{in} = 1$) digital experience to clients and their end customers. Mphasis' Service Transformation approach helps 'shrink the core' through the application of digital technologies across legacy environments within an enterprise, enabling businesses to stay ahead in a changing world. Mphasis' core reference architectures and tools, speed and innovation with domain expertise and specialization are key to building strong relationships with marquee clients. To know more, please visit www.mphasis.com

For more information, contact: marketinginfo@mphasis.com

USA
460 Park Avenue South
Suite #1101
New York, NY 10016, USA
Tel.: +1 212 686 6655

UK
88 Wood Street
London EC2V 7RS, UK
Tel.: +44 20 8528 1000

INDIA
Bagmane World Technology Center
Marathahalli Ring Road
Doddanakundhi Village
Mahadevapura
Bangalore 560 048, India
Tel.: +91 80 3352 5000



VAS 18 19 US LETTER BASL 6/72