



Handling Errors during Bulk DML operations

Erik van Roon



Who Am I?

Erik van Roon



EvROCS
COMPLETING THE PUZZLE



Member of Symposium 42

<https://sym42.org/>



Core team
MASH Program



Oracle ACE
Pro



: erik.van.roon@evrocs.nl



: www.evrocs.nl



: @evrocs_nl



Mentor and Speaker Hub

Our goal is to *connect* speakers with mentors to assist in *preparing* technical sessions and *improving* presentation skills

Interested? Read more and get in touch

<https://mashprogram.wordpress.com>



500+ technical experts helping peers globally

The **Oracle ACE Program** recognizes and rewards community members for their technical and community contributions to the Oracle community

3 membership tiers



For more details on Oracle ACE Program:
ace.oracle.com



Nominate
yourself or someone you know:

ace.oracle.com/nominate

Connect:  aceprogram_ww@oracle.com

 Facebook.com/OracleACEs

 [@oracleace](https://twitter.com/oracleace)



Background of this talk

Did a number of data migrations:

- Merge data of acquired competitor into own database
- Similar data but completely different data models
- Transformation may/will lead to problems
- Errors should be handled gracefully
- An error in a child record sometimes means the parent shouldn't be present either
- Limited window for executing the migration

In this presentation

Scripts will be mentioned like this



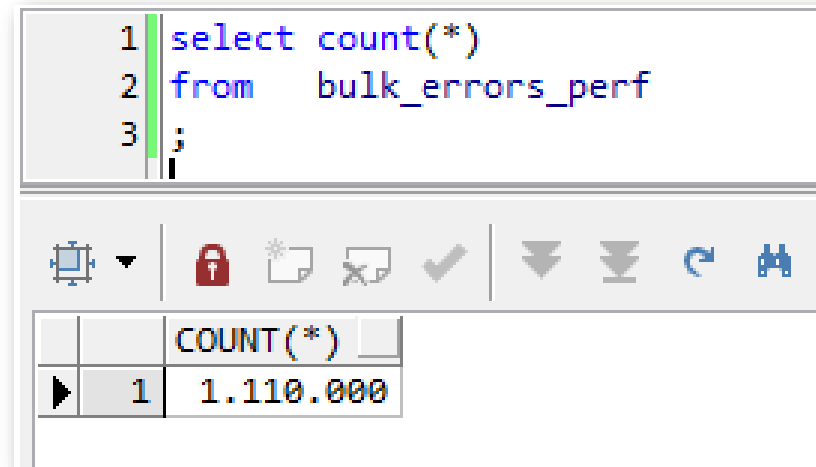
Preparation\create_demo_objects.sql

The scripts need some objects created by



Among which a table based upon sh.customers

```
1 select count(*)
2 from bulk_errors_perf
3 ;
```



	COUNT(*)
1	1.110.000

When you're done




Cleanup\cleanup.sql

So why Bulk Operations?



Row By Row



Not the fastest
But fast



And Full Control



Single SQL Statement



How fast is Bulk Fetching?

Slow-by-Slow fetching

```
open c_err;  
loop  
  fetch c_err  
  into r_err;  
  
  exit when c_err%notfound;
```

```
ERO@BULK>@Perf\slow_by_slow_fetch.sql  
Records fetched: 1.110.000  
  
PL/SQL procedure successfully completed.  
Elapsed: 00:00:05.295
```

Perf\slow_by_slow_fetch.sql

Bulk fetching

```
open c_err;  
loop  
  fetch c_err  
  bulk collect  
  into a_err  
  limit cn_bulk_limit;  
  
  exit when a_err.count = 0;
```

```
ERO@BULK>@Perf\bulk_fetch.sql  
Records fetched: 1.110.000  
  
PL/SQL procedure successfully completed.  
Elapsed: 00:00:00.709
```

Perf\bulk_fetch.sql

How fast is Bulk Data Manipulation?

Slow-by-Slow update

```
for i_err in 1 .. a_err.count
loop
  update bulk_errors_perf
  set    cust_last_name = trim(cust_last_name)
  where  rec_id          = a_err(i_err).rec_id
  ;
end loop;
```

Perf\slow_by_slow_update.sql

```
ERO@BULK>@Perf\slow_by_slow_update.sql
Runtime (sec) : 79,44
Records updated: 1.110.000
```

Bulk update

```
forall i_err in indices of a_err
  update bulk_errors_perf
  set    cust_last_name = trim(cust_last_name)
  where  rec_id          = a_err(i_err).rec_id
  ;
```

Perf\bulk_update.sql

```
ERO@BULK>@Perf\bulk_update.sql
Runtime (sec) : 12,96
Records updated: 1.110.000
```

But what if



Error Handling

During Bulk Fetching...

None

During Bulk Data Manipulation...

~~No exception handling per row~~

No savepoint to rollback to

Luckily there are Save Exceptions and Log Errors



Save Exceptions

Saves the exceptions until all iterations of the forall are processed

```
declare
  type l_tst_aat is table of some_table%rowtype index by pls_integer;
  l_tst_aa      l_tst_aat;
begin

  [.....]

  forall i_tst in 1 .. l_tst_aa.count save exceptions
    insert
    into   some_table
    values l_tst_aa(i_tst)
  ;
end;
/
```

And then....

ORA-24381: error(s) in array DML

ORA-24381

This can be handled in an exception handler

```
declare
  e_bulk_errors      exception;
  pragma exception_init(e_bulk_errors, -24381);
begin
  [...]

  forall [...] save exceptions
    [...]
;
exception
  when e_bulk_errors
  then
    [...]
end;
/
```

Pseudocollection

Pseudocollection `sql%bulk_exceptions` is available

"Composite attribute that is like an associative array of information about the DML statements that failed during the most recently run FORALL statement"

Two attributes:

Error_index

The number of the DML statement that failed

Error_code

The Oracle Database error code for the failure

Handling the exceptions

Pseudocollection identifies which statements in forall raised which exception.

```

exception
  when e_bulk_errors
  then
    dbms_output.put_line ('Exceptions');
    dbms_output.put_line ('=====');

    for i_err in 1 .. sql%bulk_exceptions.count
    loop
      l_error_index := sql%bulk_exceptions(i_err).error_index;
      l_error_code  := sql%bulk_exceptions(i_err).error_code ;

      dbms_output.put_line ('Exception sequence: ' || i_err);
      dbms_output.put_line ('Error_index      : ' || l_error_index);
      dbms_output.put_line ('Error_code     : ' || l_error_code);
      dbms_output.put_line ('Error Message  : ' || sqlerrm(-1 * l_error_code));
      dbms_output.put_line ('ID             : ' || l_val_aa(l_error_index).id);
      dbms_output.put_line ('Value         : ' || l_val_aa(l_error_index).value);
    end loop;
  
```

```

Exceptions
=====
Exception sequence: 1
Error_index      : 3
Error_code       : 2290
Error Message    : ORA-02290: check constraint (.) violated
ID               : 333
Value           : 0

Exception sequence: 2
Error_index      : 6
  
```

Notice:

- That the error code is a positive number
- How we may lose information (depending on the exception raised) because we only have the error code:
ORA-01476: divisor is equal to zero ✓
ORA-02290: check constraint (.) violated ✗

```
Exceptions
=====
Exception sequence: 1
Error_index       : 3
Error_code        : 2290
Error Message     : ORA-02290: check constraint (.) violated
ID                : 333
Value             : 0

Exception sequence: 2
Error_index       : 6
```


Sparse collections

Remember

error_index is "The number of the DML statement that failed"

For sparse collections

error_index <> index of erroneous record

error_index = x means

The x-th record in the original collection raised an exception

Not: the record at index x

Sparse Collection

Index	RecNo	Value_1	Value_2
1	1	B	C
4	2	E	F
6	3	H	I
7	4	K	L
8	5	N	O

If index 4 & 7 in this collection cause an exception

Then this is what your SQL%BULK_EXCEPTIONS will be:

Error_index	Error_code
2	1
4	1476

error_index, solutions

Loop through original collection
counting to error_index-th record

Make the sparse collection dense again

Avoid sparse collections,
mark rejected records using an extra status attribute

Solution "count"

sql%bulk_exceptions

Error_index	Error_code
2	1
4	1476
8	1
14	1
87	2290

Collection

Index	RecNo	Value_1	Value_2
1	1	B	C
4	2	E	F
6	3	H	I
7	4	K	L
8	5	N	O

Handle Exception

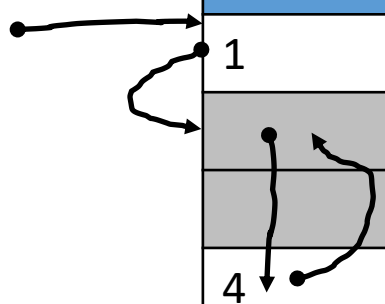
Etc.

disadvantage: extra loop through the collection
(though only if there **are** exceptions)

Solution "make dense"



Collection



A diagram on the left side of the table shows a sequence of arrows indicating the reordering of rows. It starts at the top left, points to the first row (Index 1), then to the second row (Index 4), then to the third row (Index 6), then to the fourth row (Index 7), and finally to the fifth row (Index 8). This illustrates the process of making the table dense by reordering rows based on their RecNo values.

Index	RecNo	Value_1	Value_2
1	1	B	C
4	2	E	F
6	3	H	I
7	4	K	L
8	5	N	O

Solution "make dense"



Collection


Index	RecNo	Value_1	Value_2
1	1	B	C
2	2	E	F
4	3	E	F
6	4	H	I
7	5	K	L
8	6	N	O

Solution "make dense"



Collection

Index	RecNo	Value_1	Value_2
1	1	B	C
2	2	E	F
6	3	H	I
7	4	K	L
8	5	N	O

A black arrow points from the left side of the table to the gap between the row with Index 2 and the row with Index 6, indicating the need to re-index the rows to be dense.

Solution "make dense"



Collection

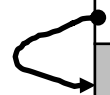
Index	RecNo	Value_1	Value_2
1	1	B	C
2	2	E	F
3	3	H	I
7	4	K	L
8	5	N	O

Solution "make dense"



Collection

Index	RecNo	Value_1	Value_2
1	1	B	C
2	2	E	F
3	3	H	I
4	4	K	L
8	5	N	O


A hand-drawn black arrow points from the index '4' to the first empty row below it, indicating the gap in the index sequence.

Solution "make dense"



Collection

Index	RecNo	Value_1	Value_2
1	1	B	C
2	2	E	F
3	3	H	I
4	4	K	L
5	5	N	O



Process this dense collection

disadvantage: extra loop through the collection
Either always, just before the forall
Or only do this, just before exception handling

Solution "use status"




Don't throw away records you don't need/want but rather set the status indicator

Index	RecNo	Value_1	Value_2	Status / Action
1	1	B	C	
2	2	BBB	CCC	
3	3	BCD	DCB	
4	4	E	F	
5	5	EEE	FFF	
6	6	H	I	
7	7	K	L	
8	8	N	O	

Solution "use status"

Don't throw away records you don't need/want but rather set the status indicator

Index	RecNo	Value_1	Value_2	Status / Action
1	1	B	C	OK / Process
2	2	BBB	CCC	NOK / Reject
3	3	BCD	DCB	NOK / Reject
4	4	E	F	OK / Process
5	5	EEE	FFF	NOK / Reject
6	6	H	I	OK / Process
7	7	K	L	OK / Process
8	8	N	O	OK / Process

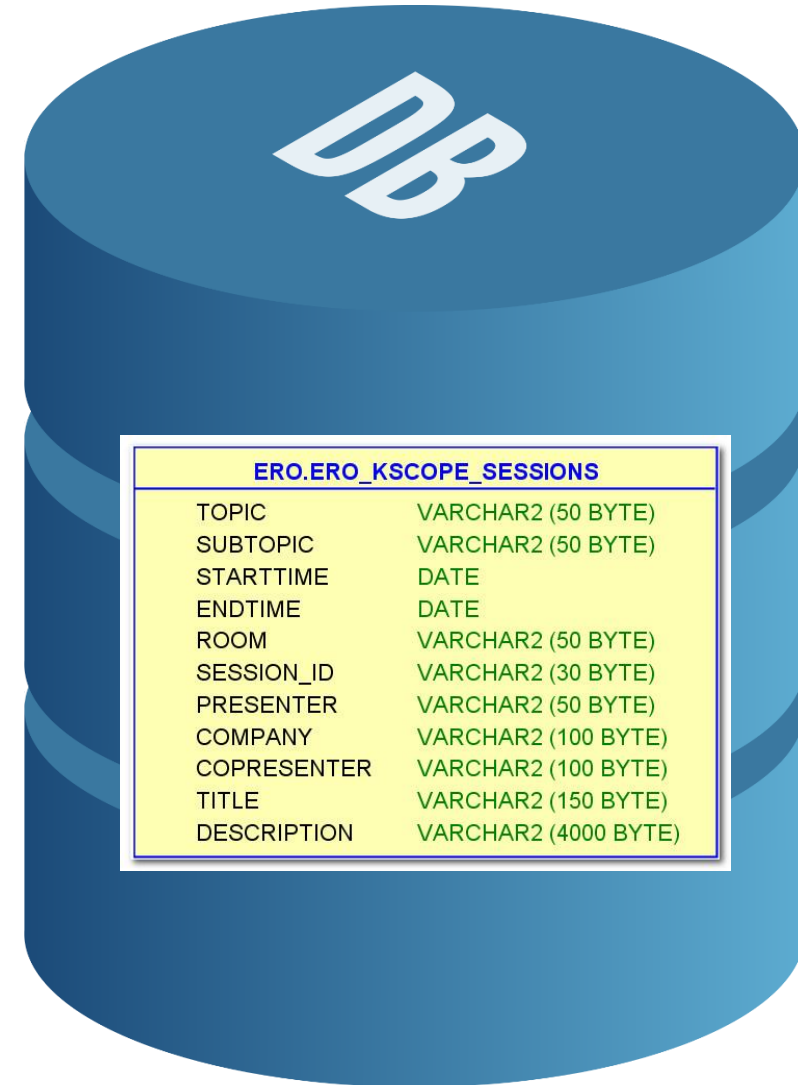


Process this dense collection, DML only for 'OK' records

disadvantage:

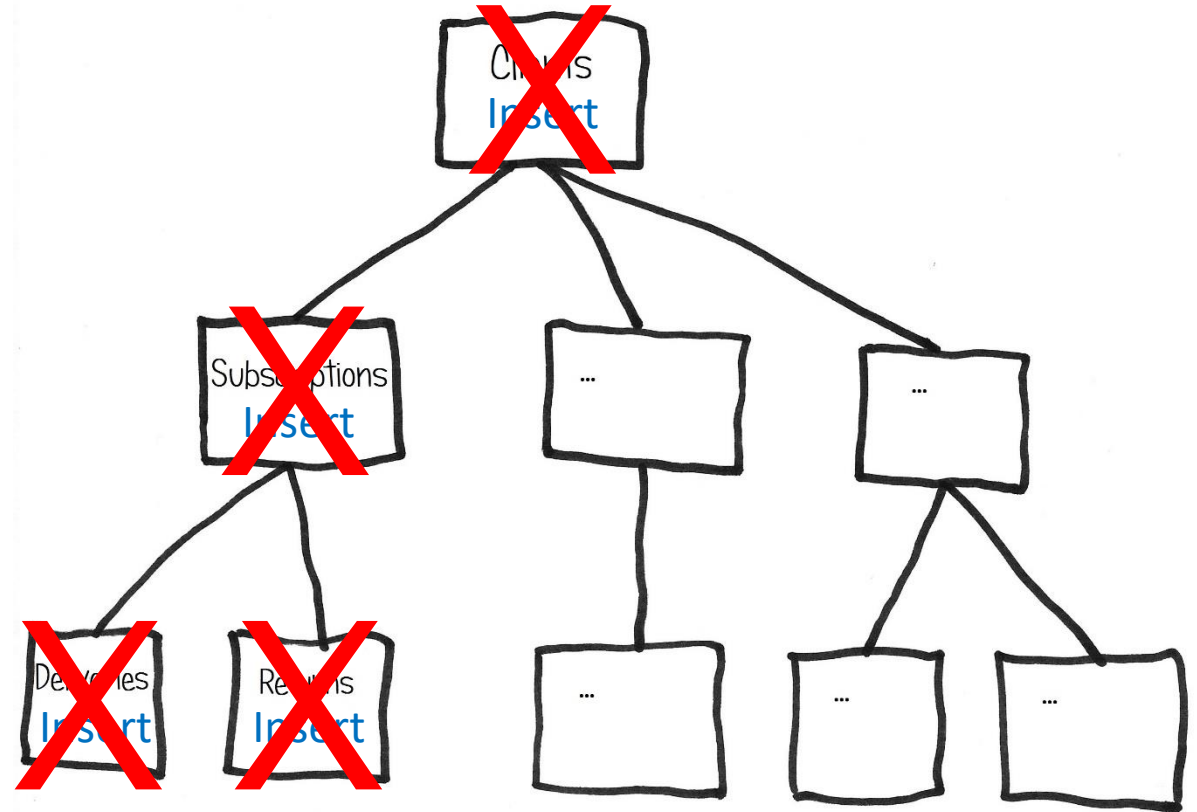
DML statements for rejected records that won't do anything (fast per record, but 'unnecessary')

But is your database....
as simple as this one?



Multiple Foralls

So, we want to prepare data for several tables and do a forall-insert for each



What if an error anywhere in the tree means the entire tree for that master record must be rejected?

Save Exceptions – Multiple Foralls

Option 1: Foreign Keys

Temporarily change the FK's to "Cascading Delete"

- Usually not an option for live databases
- Only works if the tables actually have FK's
- Logging of these 'extra' rejections require extra queries
- Requires good administration of original states of FK's and correctly reestablishing these states, whatever happens

Save Exceptions – Multiple Foralls

Option 2: Add the PK/UK of its parent to each collection

For each bulk exception:

- Determine PK/UK of the parent of the child table
- Delete all rows with same parent-PK in the child table
- If parent has other child tables, step by step delete tree below
- Delete the parent record
- Determine PK/UK of the parent of this parent
- Etc, etc, etc

Gets really ugly, really fast

Save Exceptions – Multiple Foralls

Option 3: Create a Metadata-Collection

Create an extra collection (preferably nested table)

Map each index of each collection in the tree
to the PK/UK of top-level parent

For a record that causes an exception

Get PK/UK of top-level parent from the Metadata

For each table in the tree

From the metadata determine the rows belonging to the top-level-parent

Delete every record belonging to the top-level-parent

Metadata Example

Nested table for mapping could be like:

```
create or replace type mapping_ot force as object
  (table_name          varchar2(30)
  ,collection_index   integer
  ,top_level_pk       varchar2(50)
  ,table_pk           varchar2(100)
  );
/
```

```
create or replace type mapping_ntt force as table of mapping_ot
/
```



Metadata Example

Mapping Nested Table

Tablename	Collection Index	Top Level PK	Collection PK
Clients	1	ABC	ABC
Clients	2	KLM	KLM
Subscriptions	1	ABC	S001
Subscriptions	2	ABC	S002
Subscriptions	3	ABC	S003
Subscriptions	4	KLM	S004
Deliveries	1	ABC	D101
Deliveries	2	ABC	D653
Deliveries	3	KLM	D871

If delivery at collection index 2 fails,
we can query the Nested Table to instantly find out the PK of its top-parent (ABC)

Then we can easily identify for any table the PKs that belong to the same top-parent (ABC)
by querying the Nested Table again



Metadata Example

Deleting the records for the same top-level PK/UK:

Suppose PK 'R-1534' in table RETURNS raised an exception

```
delete
from subscriptions
where subscriptions_pk in
  (select to_remove.table_pk
   from table (mapping_nt) erroneous
   join table (mapping_nt) to_remove
     on to_remove.top_level_pk = erroneous.top_level_pk
  where erroneous.table_name = 'RETURNS'
     and erroneous.table_pk = 'R-1534'
     and to_remove.table_name = 'SUBSCRIPTIONS'
  );
```

SaveExceptions\MultiForall.sql

Well, not *completely* different: LOG ERRORS

AND NOW
FOR SOMETHING
COMPLETELY
DIFFERENT

What is LOG ERRORS?

DML-statement clause

Logs errors in table with structure:

Name	Null?	Type
ORA_ERR_NUMBER\$		NUMBER
ORA_ERR_MESG\$		VARCHAR2(2000)
ORA_ERR_ROWID\$		ROWID
ORA_ERR_OPTYP\$		VARCHAR2(2)
ORA_ERR_TAG\$		VARCHAR2(2000)
[COLUMNS THAT NEED TO BE LOGGED]		VARCHAR2(4000)

Available since 10.2

Columns to be logged are all Maximum-length character datatype (or RAW)

LOG ERRORS, create log table

Creation of logtable:

Manually (obey required structure!)

Or have Oracle do it for you

```
dbms_errlog.create_error_log
(dml_table_name      in varchar2      -- table to create log-table for
, err_log_table_name in varchar2 := null -- name for log-table
, err_log_table_owner in varchar2 := null -- owner for log-table
, err_log_table_space in varchar2 := null -- tablespace for log-table
, skip_unsupported   in boolean   := false -- include unsupported datatypes?
);
```

If you're using **extended datatypes**, you may want to do it **manually**.

Having each column in the table represented in the error table by a varchar2(32767) might not be what you want.



Parameters for create_error_log

err_log_table_name

If omitted

'ERR\$_' || table_name

Truncated at maximum tablename-length (30 / 128)

If resulting tablename already exists

ORA-00955: name is already used by an existing object

Parameters for create_error_log

err_log_table_owner

If omitted: the currently connected user

err_log_table_space

If omitted: The default tablespace of err_log_table_owner

skip_unsupported

If table contains columns with unsupported datatypes

True: those columns won't be in log table

False: ORA-20069: Unsupported column type(s) found

unsupported datatypes are:

Long, *LOB, Bfile, Abstract Data Type(ADT)

How to log errors?

Add the LOG ERRORS clause to the end of the DML statement

```
log errors  
  [into logtable-name]  
  [( 'Tag' )]  
  [reject limit integer|unlimited]
```

(Optional) clauses

[**into** logtable-name]

Names the table in which logging is to be inserted

If no into clause is used, same default as with
create_error_log procedure

```
'ERR$_' || tablename
```

(Optional) clauses

[('Tag')]

Is a value for column `ORA_ERR_TAG$` in the log table

Can be used to identify the log records for this statement

If not supplied `ORA_ERR_TAG$` column will be null

(Optional) clauses

`[reject limit integer|unlimited]`

If **more than** errors than this occur, the entire statement fails

All errors are still logged into the log table

The exception raised = the last error that occurred

Default is 0, so any error will crash the statement

A statement could look like this

```
insert
into   some_table
      (id
       ,first_column
       ,second_column
       )
select some_sequence.nextval
,       some_column
,       some_other_column
from   some_other_table
log errors
into err$_some_table
('my insert')
reject limit unlimited
```

Example

LogErrors\log_errors_example.sql



Create a table

```
create table log_error_example_table
(id      integer generated by default on null as identity
,value  integer not null
);
```

Constraint says
value must be
more than 0

```
alter table log_error_example_table
add (constraint log_error_example_table_ck1
     check (value > 0)
);
```

Create a log-table

```
begin
  dbms_errlog.create_error_log
    (dml_table_name      => 'log_error_example_table'
    ,err_log_table_name  => 'log_error_example_error'
    );
end;
/
```

Example

LogErrors\log_errors_example.sql



Insert rows....

Data table contents

ID	VALUE
2	20
4	40
8	80
10	100

Log table contents

```
insert
into log_error_example_table
(value
)
select case
    when mod(level, 3) = 0 then 0
    when mod(level, 2) = 0 then level * 10
end
from dual
connect by level <= 10
log errors
into log_error_example_error
('My example of log_errors')
reject limit unlimited
;
```

ORA_ERR_NUMBERS\$	ORA_ERR_MESG\$	ORA_ERR_ROWID\$	ORA_ERR_OPTYP\$	ORA_ERR_TAG\$	ID	VALUE
1.400	ORA-01400: cannot insert NULL into ("ERO"."LOG_ERROR_E	I	My example of log_errors ...	1	...
2.290	ORA-02290: check constraint (ERO.LOG_ERROR_EXAMPLE_TAB	I	My example of log_errors ...	3	0
1.400	ORA-01400: cannot insert NULL into ("ERO"."LOG_ERROR_E	I	My example of log_errors ...	5	...
2.290	ORA-02290: check constraint (ERO.LOG_ERROR_EXAMPLE_TAB	I	My example of log_errors ...	6	0
1.400	ORA-01400: cannot insert NULL into ("ERO"."LOG_ERROR_E	I	My example of log_errors ...	7	...
2.290	ORA-02290: check constraint (ERO.LOG_ERROR_EXAMPLE_TAB	I	My example of log_errors ...	9	0

Extra information in the log table...

```
alter table log_error_example_error
add (dml_timestamp          timestamp          default systimestamp
     ,dml_sessionid        integer            default sys_context('userenv', 'sessionid')
     ,dml_user              varchar2(128)       default sys_context('userenv', 'current_user')
     ,dml_terminal          varchar2(128)       default sys_context('userenv', 'terminal')
     ,dml_os_user           varchar2(128)       default sys_context('userenv', 'os_user')
     ,dml_nls_date_format   varchar2(128)       default sys_context('userenv', 'nls_date_format')
     ,dml_nls_date_language varchar2(128)       default sys_context('userenv', 'nls_date_language')
);
```

Running the same insert as before....

ERR_TAG\$	ID	VALUE	DML_TIMESTAMP	DML_SESSIONID	DML_USER	DML_TERMINAL	DML_OS_USER	DML_NLS_DATE_FORMAT	DML_NLS_DATE_LANGUAGE
sample of log_errors	1		29-MAY-22 02.14.05.570913 PM	1.420.001	ERO	LAPTOP-EVROCS	Erik	DD-MON-RR	AMERICAN
sample of log_errors	3	0	29-MAY-22 02.14.05.571433 PM	1.420.001	ERO	LAPTOP-EVROCS	Erik	DD-MON-RR	AMERICAN
sample of log_errors	5		29-MAY-22 02.14.05.571522 PM	1.420.001	ERO	LAPTOP-EVROCS	Erik	DD-MON-RR	AMERICAN
sample of log_errors	6	0	29-MAY-22 02.14.05.571685 PM	1.420.001	ERO	LAPTOP-EVROCS	Erik	DD-MON-RR	AMERICAN
sample of log_errors	7		29-MAY-22 02.14.05.571792 PM	1.420.001	ERO	LAPTOP-EVROCS	Erik	DD-MON-RR	AMERICAN
sample of log_errors	9	0	29-MAY-22 02.14.05.571923 PM	1.420.001	ERO	LAPTOP-EVROCS	Erik	DD-MON-RR	AMERICAN

Some advantages over save exceptions are

- Works for each and every DML statement, not just FORALL
- We have the entire error message, not just error number
- We (can) have every value of every column as it arrives at the table, not just the ones we have in our statement (for example, also the values supplied by triggers and defaults)

Some things to be aware of

- No automatic clean up
Table grows over time
'Old' errors may contaminate your query if not sensibly tagged
- Multiple error tables on single table possible
Can be just what you want/need.
Can also lead to developers all creating their own error table
- The errorlog table is in no way 'connected' to the data table
So dropping or altering one has no effect on the other

Some things to be aware of

- Identification of errors of 'last DML performed' only by TAG
So sensible tagging is essential to retrieve the correct errors
- Log table is not session specific
Unless it's a Global Temporary Table (see later)
Or you add the information yourself
- Structure of log table needs to be kept in sync with original
No errors if not, but you may lose information (see later)
- The logging is committed inside an autonomous transaction
- Statement will always complete successfully (if within reject limit).
To know if records were rejected you need to query the Log table.

Log table as Global Temporary Table (GTT)

- Must be created manually
- Must be defined as "on commit preserve" or the commit in the autonomous transaction will remove your logging

Possible advantages to a GTT log table

- All visible logging is done by 'your' session
So TAG only needs to identify statements in this session
- Cleanup automatically at end of session
- Table can be truncated during a session without hurting other sessions

Disadvantage to a GTT log table

- If error logging needs to be preserved, manually insert it into another table is needed

Identification based on TAG is tricky

- Procedure name is going to be the same in the next call
- Session id's are being reused
- Date-time strings may not be unique in multiuser environment

A possible solution

- Use SYS_GUID to generate a Globally Unique Identifier at the start of a transaction and use it to identify its logged errors

What if structure of original table differs from log table?

- **Column added to or dropped from original table**

No problem, only columns present in both tables are logged

After also adding a new column to the log table it's also logged

- **Order of columns differ**

Again, no problem. Columns are logged as expected

- **Column in log table smaller than data in DML statement**

Statement crashes with

ORA-38906: insert into DML Error Logging table "[.]" failed

followed by the error message that was attempted to be logged


Multiple statements

What if we have the same situation we discussed for forall:
Several tables, and if anywhere in the tree a record fails, the whole tree for the top-level record it belongs to must be rejected?

Multiple statements

We could:

- Get the FK of erroneous records from errorlog table
- Get the PK of the top-level record
- Based on the top-level PK delete the entire tree



Performance

(Stolen from Oracle-Base, Tim Hall)

Test of insert of 100.000 rows with 2 errors on different database version (on different servers):

	10.2.0.4	11.2.0.3	11.2.0.4	12.1.0.1
	=====	=====	=====	=====
DML Error Logging :	07.62	08.61	04.82	00.94
DML Error Logging (APPEND) :	00.86	00.38	00.85	01.07
FORALL ... SAVE EXCEPTIONS :	01.15	01.01	00.94	01.37

times are in seconds

Summary

- Log Errors gives the actual error message
Save Exceptions only gives the error code
- Log Errors also stores the actual data
Save Exceptions only has pointers (which are tricky for sparse collections) to data in another collection
- Errors Logged by Log Errors is persistent
Save Exceptions errors are volatile
- Log Errors can be used for the ultimate bulk operation:
a single DML statement
Save exceptions can only be used for Forall statements

Summary

- With Save Exceptions, correcting executed DML in a tree of tables requires extra collections and extra plsql coding.
With Log Errors Just some extra DML statements are needed
- With Log Errors the statement always succeeds
(when errors \leq reject limit)
Save Exceptions raises an exception that can be handled
- Log Errors makes you work to identify the errors caused by the last statement
Save exceptions only has the errors for the last statement
- Before 12c performance of Log Errors may be bad



"Stupid questions do exist.

But it takes a lot more time and energy to correct a stupid mistake than it takes to answer a stupid question, so please ask your stupid questions."

a wise teacher who taught me more than just physics

Thanks !!!