



An Alternative Approach to Summary Table Management Using Advanced Queues

By Arup Nanda

One of the biggest challenges facing the designer of a data warehouse demanding near real time synchronization is that of maintaining summary tables. While it becomes imperative to incrementally refresh the changes from the source tables, traditional methods such as materialized views fail when the summary building query becomes extremely complex in nature with several joins, especially outer joins. The complete refresh may not be acceptable due to the real time performance needs. Using triggers to populate the summary table does not work either, due to the data concurrency requirements. This article presents a rather unconventional approach to resolve this issue using Oracle Advanced Queues. The solution is to feed the changes to the source tables from multiple processors to a queue and at the other end, have only one process retrieve the information and apply the changes to the summary tables. This approach takes care of the concurrency problem while attaining a near real time performance.

Background

In the actual case shown here, we had a very complex query that the client applications issued before processing other records. Note that the tables and columns have been changed to protect the intellectual property.

```
SELECT DISTINCT NVL(C.COL1, R.COL1) AS COL1,
RQ.COL2, COUNT(C.COL3) AS COUNTER
FROM TAB1 R, TAB2 C, TAB3 CP, TAB4 RQ
WHERE R.COL4 = C.COL4
AND C.COL1 = RQ.COL1
AND ( CP.COL5 = :b0
OR (CP.COL5 = :b1 AND CP.COL6 = :b2))
AND CP.COL3 = C.COL3
AND NOT EXISTS
(SELECT 1
FROM TAB5 CCL,TAB6 CL, TAB7 CHO
WHERE CCL.COL7 = CL.COL7
AND CL.COL8 = CHO.COL8
AND CCL.COL3 = C.COL3
AND CHO.COL1 = 'Y')
GROUP BY NVL(C.COL1, R.COL1), RQ.COL2
ORDER BY RQ.COL2
```

The initial design was accomplished by creating a view from the query and letting the users select from the view. This proved costly in terms of performance. Running the query in SQL*PLUS took about forty minutes to run, and also took up huge temporary tablespace areas. Since this query was supposed to be run very frequently, this approach could not be used.

The second approach was to use a materialized view refresh-on-commit for the query. Since the query was complex, it could not be set up for fast refresh. A complete refresh was taking more than 24 hours and thus was not feasible. Other options considered along this line were building materialized views from materialized views and making each one fast refreshable on commit. However, the maintenance of several layers of materialized views would have been expensive in terms of personnel resources and space. A bigger problem in on-commit fast refreshable materialized views was the result of a failure. If a refresh ever failed, the future fast refreshes would be stopped until the DBA made a call to DBMS_MVIEW package. Since this type of refresh has a high transaction rate it was prone to errors. The overall solution would have proved too expensive to maintain. In addition, when a prototype was built and tested at the expected transaction rate, the ultimately materialized view could not be refreshed within an acceptable time limit.

Ultimately the only acceptable option was to create triggers on all source tables to capture changes and update the summary table built earlier. Since the query's nature was known, it was easy to assemble the code that updated the counts in the summary table when a record from source tables was changed in such way that its selectivity in the query was also altered.

Space constraints do not permit a complete presentation of the actual solution for the case given above. Therefore, it will be illustrated using a much-simplified query based on the sample tables from the SCOTT schema as shown here:

```
SELECT DEPTNO, COUNT(*)
FROM EMP
WHERE STATUS = 'ACTIVE'
GROUP BY DEPTNO
```

Although it was not a complex query, it will suffice for the demonstration. Since this query was taking a long time to execute, we built a table called DEPT_COUNTS with two columns, DEPTNO and EMP_COUNT. The table was

continued on page 25

initially populated by a query. A trigger was created on the EMP table that updated the EMP_COUNT when the rows were changed, deleted or inserted. The pseudo code was something like the following.

```

If INSERTING then
  If STATUS = 'ACTIVE' then
    If (record present in DEPT_COUNTS for this deptno) then
      Update DEPT_COUNTS
      Set EMP_COUNTS = EMP_COUNTS + 1;
    Else
      Insert into DEPT_COUNTS values
      (this Deptno, 1);
    end if
  end if;
...
    
```

The code was then extended for deletions and updates in the same manner. Since this trigger manipulated the counts, the needed data could be obtained from the DEPT_COUNTS table without using the complex query.

This approach was fine in principle and would have worked in a single user environment. However, if two sessions insert employee records for the same department, for which there were no records before, then the following situation arises.

Time	Action
0	There is no employee record with DEPTNO = 5; therefore there is no record in DEPT_COUNTS for DEPTNO = 5
1	Session 1 inserts record with DEPTNO = 5; does not commit.
2	The trigger finds no record for DEPTNO = 5 in the DEPT_COUNTS table. This inserts a record for DEPTNO = 5 in DEPT_COUNTS table.
3	Session 2 inserts another record with DEPTNO = 5
4	The trigger sees the read consistent view of the table DEPT_COUNTS. In this read consistent view, it also finds no record for DEPTNO = 5 in the DEPT_COUNTS table. This inserts a record for DEPTNO = 5 in DEPT_COUNTS table. However, since the primary key is DEPTNO, and for DEPTNO = 5 there is another row, albeit uncommitted, this transaction waits for a lock.
5	Session 1 Commits. The committed value of EMP_COUNTS = 1.
6	Session 2 errors out, since it also tried to insert the same record, DEPTNO = 5 and EMP_COUNTS = 1.

As you can see the second transaction fails when it should not have. A similar problem occurs when a session deletes the last employee of a department but the other session does not see that and tries to update it. Yet another problem is locking. If two sessions enter records for the same DEPTNO, the second one must wait to get the lock until the first one commits. This creates an artificial row locking conflict.

All these significant problems required coming up with another solution. At this point, the advanced queue approach was considered. Advanced queues are typically used for application development where one part of the processing is de-coupled from the others. However, the model may be used in various other situations such as this one. In order to present the concepts as concisely as possible, the same simplified example as above will be used. Although this scenario is much less complicated than the actual case, the design of the solution is still useful.

What are Advanced Queues (AQs)?

Advanced Queue (AQ) is a queue-based system entirely within the database. A queue is a simple concept to grasp. It is like a pipeline where anyone can place something in a queue where it waits until it is selected. This is typically done on a first in/first out basis.

In order to explain advanced queues more easily, some definitions are needed:

- **Producer:** Person placing information in the queue
- **Message:** Information that goes into the queue
- **Consumer:** Person who selects the message from the queue
- **Payload:** Message carrying the useful data. A message is like an envelope enclosing the payload (letter).
- **Enqueuing:** The act of placing a payload in the queue
- **Dequeuing:** Retrieving the payload from the queue
- **Subscribers:** Processes that get the message off of the queue

There can be more than one producer and consumer for a queue. The producer and consumers can also be the same process. A queue can have subscribers, but this does not prevent other subscribers from dequeuing the same message. This is called a multi-consumer publishers-subscriber model queue. There can be rules defined that allow certain, but not all, subscribers to dequeue payloads from the queue.

AQs can be persistent or non-persistent. The former type is not in memory but physically stored. In case of a database shutdown, the queue contents are preserved. This type of queue is also protected by the database; it can be backed up and recovered. The latter type of queue is in memory only and disappears as soon as the instance is brought down. Needless to say, the performance is better in non-persistent queues.

Perhaps the single most important attribute of an AQ is its ability to participate in a transaction – i.e. the message is placed in the queue only when the overall transaction commits. It is never placed if the transaction is rolled back. This characteristic of AQ provides rich benefits for building a messaging system that is transaction-aware or, in other words, able to handle real-world business processes.

In effect, an AQ-based system is like a huge delivery system. Since it is contained entirely within an Oracle database, it is protected from failures and is transaction-aware. Each queue is associated with a database table called Queue Table (QT) that can hold more than one queue. Each QT can handle only one type of payload. The payload can be a simple VARCHAR2 string or an Oracle Abstract Datatype (ADT) object. Since the QT is a physical table, it can also have storage attributes such as tablespace, initial extent, etc. defined.

A queue is created in a Queue Table. A queue called “exception queue” is always created automatically in a queue table. This is used to retain the messages that failed while dequeuing. The queue system is managed by two Oracle-supplied packages named DBMS_AQ and DBMS_AQADM, in addition to the Oracle Enterprise Manager Console.

Solution Approach

Returning to the problem at hand, the proposed solution involves building a queue with the triggers as producers that place the changed data into the queue. On the other side only one consumer (a stored procedure) dequeues the message and updates the summary table. Since the update is done by only one session, the problem of incorrect updates will never occur.

Initial Database Preparation

The database has to be set up with the following parameter and value in the `init.ora` file.

```
aq_tm_processes = 2
```

This sets up two AQ Time Manager processes to check the queues for messages for time related events such as expiration, retry, etc. These can also be set up by `ALTER SYSTEM`. This parameter can have a value up to 10. The processes are started as background processes and named with a format `QMn`, where `n` is the number of the monitor.

The users who will operate the queues need to be granted appropriate privileges using the following two lines, (provided the user is `SCOTT`.) Connecting as `SYS`, issue these two statements.

```
GRANT EXECUTE ON DBMS_AQADM TO SCOTT;
GRANT EXECUTE ON DBMS_AQ TO SCOTT;
```

Setting Up the Working Procedures

The first task is to define a type that will be used in the change data transfer. This type is an Oracle Abstract Data Type (ADT) and will be used to define the payload of the message. The type is defined as follows. All statements from now on will be executed as the user `SCOTT`.

```
create or replace type dept_counts_type as object
(
  action_type  char(1),
  old_deptno   number(2),
  new_deptno   number(2)
)
```

Since we have to transmit the information about what happened and which deptno was affected, we have defined the old and the new department numbers and the action code which will be one character string with values `I`, `D` or `U`, designating Insert, Delete or Update respectively. Next, we will build the procedure that will perform the update on the summary table. One special consideration must be made concerning this operation. Since the table is maintained by us and not by Oracle, what happens when someone truncates the `EMP` table, or even drops it? These events will not fire the trigger that captures the changes. Therefore, we need to introduce another event that truncates the `DEPT_COUNTS` tables when the `EMP` table is truncated or dropped. This can be done by passing a special action code, `T` to the processing procedure as shown here:

```
create or replace procedure process_dept_counts
( p_action      IN char,
  p_old_deptno  IN number,
  p_new_deptno  IN number)
as
  invalid_action exception;
begin
  if (p_action = 'T') then
    execute immediate 'truncate table dept_counts';
  elsif (p_action = 'D') then
    update dept_counts
    set emp_count = emp_count - 1
    where deptno = p_old_deptno;
  elsif (p_action = 'I') then
    update dept_counts
    set emp_count = emp_count + 1
    where deptno = p_new_deptno;
```

```
    if (SQL%NOTFOUND) then
      insert into dept_counts
      values
      (p_new_deptno, 1);
    end if;
  elsif (p_action = 'U') then
    update dept_counts
    set emp_count = emp_count - 1
    where deptno = p_old_deptno;
    if (SQL%NOTFOUND) then
      null;
    end if;
    update dept_counts
    set emp_count = emp_count + 1
    where deptno = p_new_deptno;
    if (SQL%NOTFOUND) then
      insert into dept_counts      values
      (p_new_deptno, 1);
    end if;
  else
    raise invalid_action;
  end if;
end;
```

NOTE: To conserve space, I have not taken into account all possibilities while writing this code. For example, what happens when the `EMP_COUNT` becomes 0 in the `DEPT_COUNTS` table? The row should then be deleted. Similarly, other housekeeping tasks such as exception handling have not been considered here. The example is presented to convey the general idea; the reader is expected to consider all aspects of the system when developing for a real system.

Setting Up the Queue System

Now we will build the queue table, called `dept_counts_qt` using the code shown here:

```
begin
  DBMS_AQADM.CREATE_QUEUE_TABLE (
    queue_table => 'DEPT_COUNTS_QT',
    queue_payload_type=> 'DEPT_COUNTS_TYPE',
    multiple_consumers=>FALSE,
    storage_clause =>
    'TABLESPACE USR INITRANS 10 STORAGE (FREELISTS 10 FREELIST GROUPS 2)',
    compatible => '8.1');
end;
```

In this queue table, we will create the queue called `dept_counts_q`.

```
begin
  DBMS_AQADM.CREATE_QUEUE (
    queue_name => 'DEPT_COUNTS_Q',
    queue_table=> 'DEPT_COUNTS_QT',
    max_retries=>'5',
    retry_delay=>'0');
  /* Now Start the Queue */
  dbms_aqadm.start_queue('DEPT_COUNTS_Q',TRUE,TRUE);
end;
```

At this point, we are ready to place the messages in the queue. We will create a generic procedure to enqueue `dept_counts_q`, so that it can be used by other applications as well using the following procedure:

continued on page 27

```

create or replace procedure enq_dept_counts_q
(p_msg in dept_counts_type)
as
  enq_opt dbms_aq.enqueue_options_t;
  msg_prop dbms_aq.message_properties_t;
  msg_id raw(16);
begin
  sys.dbms_aq.enqueue (
    'DEPT_COUNTS_Q',
    enq_opt,
    msg_prop,
    p_msg,
    msg_id);
end;

```

Similarly to dequeue the payload from the queue, use another generic procedure shown here:

```

create or replace procedure deq_dept_counts_q
as
  deq_opt dbms_aq.dequeue_options_t;
  msg_prop dbms_aq.message_properties_t;
  payload dept_counts_type;
  msgid raw(16);
begin
  loop
    deq_opt.wait := dbms_aq.forever;
    deq_opt.navigation := dbms_aq.next_message;
    dbms_aq.dequeue(
      'DEPT_COUNTS_Q',
      deq_opt,
      msg_prop,
      payload,
      msgid
    );
    process_dept_counts(
      payload.action_type,
      payload.old_deptno,
      payload.new_deptno);
    commit;
  end loop;
end;

```

The final action is to send the messages to the queue via a trigger. We can create a trigger on the EMP table using the logic discussed above:

```

create or replace trigger tr_ar_iud_emp
after insert or delete or update on emp
for each row
declare
  l_action char(1);
  l_old_deptno number(2);
  l_new_deptno number(2);
begin
  if (inserting) then
    l_action := 'I';
    l_old_deptno := null;
    l_new_deptno := :new.deptno;
  elsif (deleting) then
    l_action := 'D';
    l_old_deptno := :old.deptno;
    l_new_deptno := null;
  elsif (updating) then
    l_action := 'U';
    l_old_deptno := :old.deptno;
    l_new_deptno := :new.deptno;
  else
    null;
  end if;
  enq_dept_counts_q(
    dept_counts_type(
      l_action,
      l_old_deptno,
      l_new_deptno));
end;

```

For the truncate or drop part, we need to code that logic in a database trigger as shown here:

```

create or replace trigger tr_trunc_or_drop_table
after drop or truncate on schema
declare
  ev varchar2(200);
  ow varchar2(200);
  ty varchar2(200);
  ob varchar2(200);
  errm varchar2(2000);
  l_msg dept_counts_type :=
    dept_counts_type('T',null,null);
begin
  ow := ora_dict_obj_owner;
  ty := ora_dict_obj_type;
  ob := ora_dict_obj_name;
  if ( (ow = user) and
    (ty = 'TABLE') and
    (ob = 'EMP'))
  then
    enq_dept_counts_q (l_msg);
  end if;
end;

```

Now, we have a complete system to maintain the summary table DEPT_COUNTS from the EMP table. The procedure deq_dept_counts_q is run from the SQL prompt or in the background via a shell script. It runs in an infinite loop until it is killed. The procedure continues to listen to the queue and processes the messages that pass through it. This idea can be extended to cover more than one source table and even more than one summary table. The queues can be extended to cover a variety of changes or a separate queue can be created for each type of payload. The former is needed when

the changes need to be applied as one atomic transaction. To facilitate this process, the dequeue process can be kicked off from a database job through `dbms_job`, in which case the job will automatically start during database startup, making administration even simpler.

The completely assembled system is depicted in Figure 1.

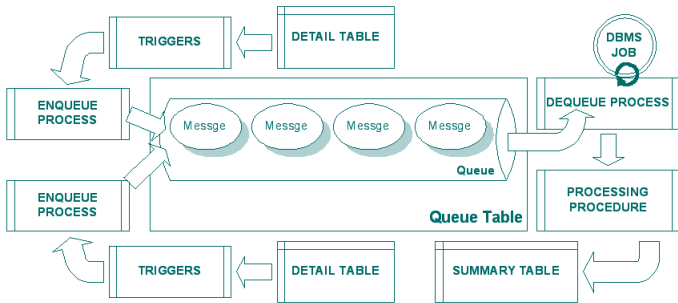


Figure 1: Diagram of completed system

Important Consideration

The whole system described in this article is based on an asynchronous and decoupled transaction-processing model. Please note that since they are decoupled, the normal database properties such as read consistency are not available. This may render the system unfeasible in specific situations. For instance, in this example, if the source session inserts into EMP table but does not commit, it will not be able to see that record in DEPT_COUNTS table. This is unlike a regular transaction processing system where the source will be able to see this record but not others. This important property must be carefully considered when applying the model to a real world system.

Administration

This system itself does not need any type of administration; however sometimes it may be worthwhile to “peek under the hood.” The frequently requested procedure is to identify the number of messages in the queue at a specific time. This can be done using the following query.

```
select count(*) from AQ$DEPT_COUNTS_QT where queue = 'DEPT_COUNTS_Q'
```

If the messages are not retrieved from the queue, check the exception queue, named after the queue table and the suffix `_QE`. If the queue table name is `DEPT_COUNTS_QT`, then exception queue is named `AQ$DEPT_COUNTS_QT_QE`.

Other Uses

As seen from the example, this AQ based system can be used to maintain any type of summary tables, not just in data warehousing environments but in practically any type of application: OLTP, Web-based, or a mixture of DSS and OLTP. The near real-time performance of this system makes it ideal for several uses without placing undue burden on the system resources as well as making users happy with faster response time. The faster response (seconds as opposed to hours) opens up another possible use of this system, namely normal applications replacing complicated queries. For instance, a banking application might use this while displaying a customized message to the customer depending upon his or her various relationships with the bank. Typically, this needs complicated queries and response time sometimes wanes. Instead of adding additional indexes which may not be feasible, the AQ system may be used to maintain a near real time summary table for the relationships, and the application can select from the table.

This example was intended to provide the concepts and workings of a functional AQ based system. Real life situations are much more complex but the inner workings remain the same. The objective of the exercise was to introduce a relatively unknown usage of an Oracle tool to solve a well-known problem.



About the Author

Arup Nanda has been an Oracle DBA for more than nine years working in all aspects of database design and management such as performance tuning and disaster recovery planning. He is the founder of Proligence, Inc. (www.proligence.com), a New York area based Oracle database consultancy firm providing specialized services such as setting up replication, parallel server, performance tuning and providing alternative solutions like the one described in this article. He can be reached at arup@proligence.com.