

# *Neat Things to Help Manage Partitioned Objects*

NYOUG June 8, 2004

# Agenda

- General issues
- Monitoring partitioned indexes
- DBMS\_STATS is a must
- Partitioning using materialized views
- List partition maintenance—the unknown frontier



# Ripple affect



- Modifications required to properly detect problems
- Partition maintenance can invalidate index partitions/sub-partitions
- Object re-compilation required as part of maintenance
  - impact scheduling
  - communication required



# Self-inflicted

Enterprise -- online



# Invalidating indexes

- Aborted direct path Loader sessions
- Direct path insert ... /\*+ append \*/
- Partition splitting when rows are in container being split
- Partition dropping when dropped partition contains rows
- Moving table partitions



6

# Monitoring Partitioned Indexes

- PARTITIONED = 'YES' in *dba\_indexes*
  - STATUS = 'N/A'
- PARTITIONED = 'NO' in *dba\_indexes*
  - STATUS = 'VALID'
- COMPOSITE = 'YES' in *dba\_ind\_partitions*
  - STATUS = 'N/A'
- COMPOSITE = 'NO' in *dba\_ind\_partitions*
  - STATUS = 'USABLE'



# Analyze



# DBMS\_STATS Highlights

- granularity—determines the type of statistics to gather for partitioned objects
- estimate\_percent—user defined number or dynamic table-specific choice made by Oracle
- degree—parallelism
- dba\_tab\_modifications—held in memory and only forced to disk when DBMS\_STATS runs





# DBMS\_STATS

## GRANULARITY

|                     |                                     |
|---------------------|-------------------------------------|
| <i>DEFAULT</i>      | Global and partition                |
| <i>SUBPARTITION</i> | Only subpartition                   |
| <i>PARTITION</i>    | Only partition                      |
| <i>GLOBAL</i>       | Global                              |
| <i>ALL</i>          | Global, partition, and subpartition |



# DBMS\_STATS

## ESTIMATE\_PERCENT

|                         |  |
|-------------------------|--|
| <i>Hard-coded</i>       | Number from 1 to 100                     |
| <i>AUTO_SAMPLE_SIZE</i> | Decision made based on condition of data |

*estimate\_percent=>dbms\_stats.auto\_sample\_size*



# DBMS\_STATS

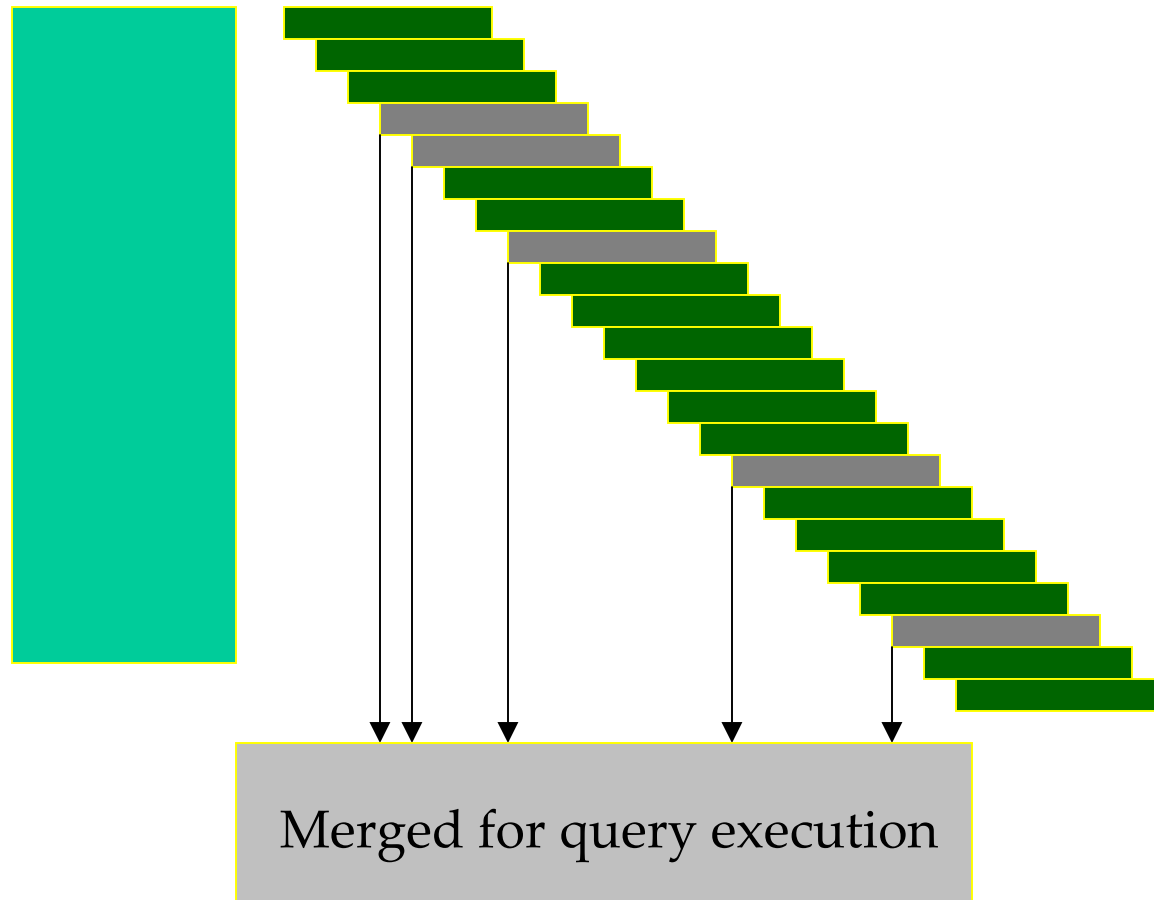
## DEGREE

|                       |                                     |
|-----------------------|-------------------------------------|
| <i>Hard-coded</i>     | Integer                             |
| <i>Null</i>           | Uses table default                  |
| <i>DEFAULT_DEGREE</i> | Influenced by initialization values |


*degree=>dbms\_stats.default\_degree*



# Global / Partition-level




# Global / Partition-level



**Analyze**

**DBMS\_STATS**



# Switch to partitioned

- No limit on size of table
  - number of rows
  - Mb, Gb, Tb
- Fat tables with fewer rows
- Real time with close to zero down time
- Application quiesced to accomplish switchover
- Keep non-partitioned as long as possible

Space permitting



# Switch to partitioned

- Make partitioned replica of non-partitioned
- Place partitioned table in nologging mode
- Place a snapshot log on non-partitioned table
- Lay snapshot on top of *prebuilt* partitioned table
- Complete refresh on snapshot





# Switch to partitioned

- Build primary key constraint on partitioned table
  - preferably local index
  - strict implementation restrictions
- Create a refresh mechanism via the job queue
- Negotiate switchover time



# The switchover

- Disable application access
- Monitor row count in snapshot log
- When zero, do the following
  - drop refresh job
  - drop snapshot
  - no rows export on non-partitioned table
  - name switching
  - build secondary indexes
  - import with ignore=y



# The switchover—caveats

- Dropping the refresh job
  - check *dba\_jobs\_running* beforehand
  - */\*+ rule \*/* hint with *9i* and *10g*
- Building secondary indexes
  - names may clash with non-partitioned index names
  - perform in parallel
  - may require partition renames afterwards



# The unknown frontier

```
SQL> create table voter (  
2     name varchar2(30),  
3     province varchar2(2))  
4     partition by list (province)  
5     (partition east values ('NF', 'PE', 'NS', 'NB'),  
6     partition quebec_ontario values ('QC', 'ON'),  
7     partition west values ('MB', 'SK', 'AB', 'BC'),  
8     partition rest values (DEFAULT));
```

Table created.



# List partitioning

- Values stored in HIGH\_VALUE in USER\_TAB\_PARTITIONS
- REST partition is the equivalent of MAXVALUE'd partition in range-based partitioning
- REST value is pseudo-boolean DEFAULT



# List partition splitting

```
SQL> alter table voter split partition quebec_ontario
2   values ('QC')
3   into
4   (partition quebec,partition ontario);
```

Table altered.

- 'QC' values to store in new partition
- **partition quebec** is the new partition created by the split



# Neat Things to Help Manage Partitioned Objects

[abbey@pythian.com](mailto:abbey@pythian.com)

