

ORACLE®

Automatic Data Optimization for Information Lifecycle Management

Hariharan Lakshmanan
Software Development Manager

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Agenda

- Information Lifecycle Management challenges
- Heat map deep dive
- ADO deep dive
- ADO implementation pointers

ILM challenges

- Manage increasing data volumes
 - Without hurting performance
 - Without growing cost
 - With minimal intervention

Data life cycle



- Recently inserted, actively updated

- Frequently Queried for Reporting

- Infrequently accessed for queries

- Retained for long term analytics and compliance with corporate policies and regulations

ILM strategy

- Data in the appropriate format based on usage patterns
- Data in the appropriate storage tier

ILM strategy continued



- Recently inserted, actively updated (advanced compression)

- Frequently Queried for Reporting (query high)

- Infrequently accessed for queries (archive low)

- Retained for long term analytics and compliance with corporate policies and regulations (archive high)

Automatic Data Optimization

Simple Declarative SQL extension



ALTER TABLE sales ILM add policy	
<ul style="list-style-type: none">▪ Advanced Row Compression (2-4x)▪ Affects ONLY candidate rows▪ Cached in DRAM & FLASH	row store compress advanced row after 2 days of no modification
<ul style="list-style-type: none">▪ Warehouse Compression(10x)▪ High Performance Storage	column store compress for query high after 1 week of no modification
<ul style="list-style-type: none">▪ Warehouse Compression(10x)▪ Low Cost Storage	tier to low cost tablespace
<ul style="list-style-type: none">▪ Archive Compression(15-50X)	column store compress for archive high after 6 months of no modification

Heat Map

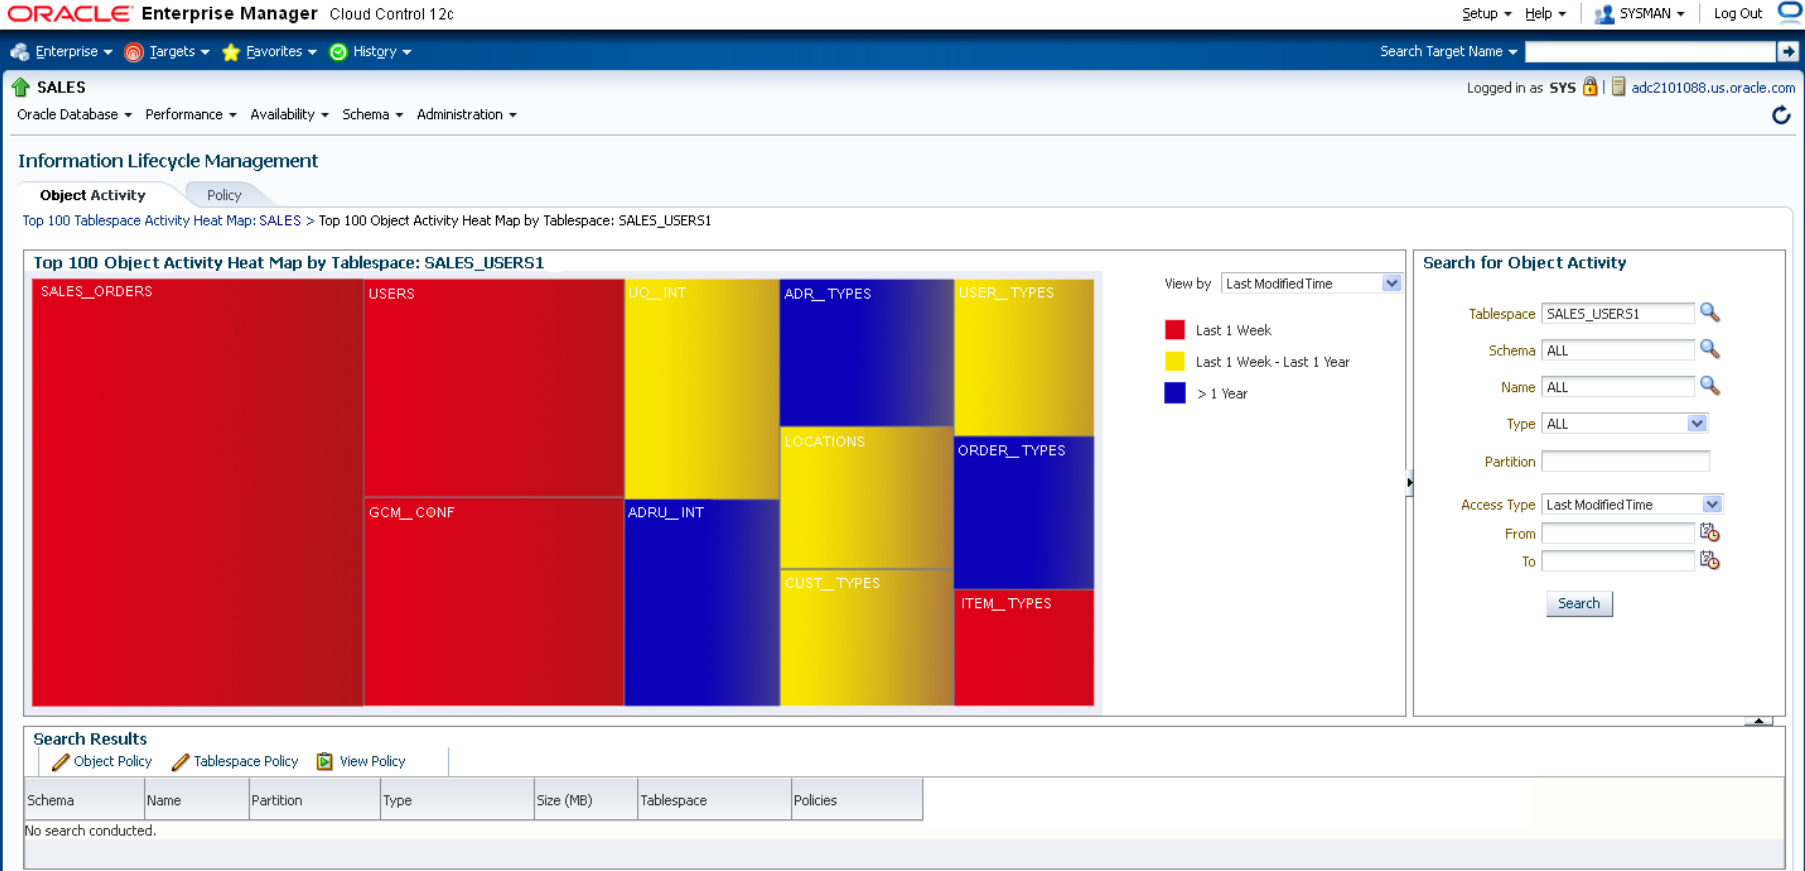
Usage Tracking



- **“Heat Map” tracking**
 - Query and modification times tracked by segment
 - Modification times tracked for database blocks
- **Comprehensive**
 - Distinguishes index lookups from full table scans
 - Automatically excludes maintenance tasks:
 - Stats, DDLs, backups, table redefinitions, etc.
- **High Performance**
 - Object level at no cost
 - Block level << 5% cost

Heat Map

Enterprise Manager Visualization



Segment level heat map views

- (USER/DBA) _HEAT_MAP_SEGMENT
- (USER/DBA) _HEAT_MAP_SEG_HISTOGRAM
- DBMS_ILM_ADMIN for setting heat map stats (testing)

Example

```
create table scott.test (empn number) tablespace tbs_1;
```

```
insert into scott.test values (1);
```

```
commit;
```

```
select owner, object_name, segment_write_time mod_time from  
  dba_heat_map_segment where owner = 'SCOTT' and object_name = 'TEST';
```

OWNER	OBJECT_NAME	MOD_TIME
SCOTT	TEST	13-MAY-2018

Example continued (dba_heat_map_seg_histogram)

```
select owner, object_name, track_time, segment_write DML,  
full_scan scan, lookup_scan idx from  
dba_heat_map_seg_histogram where owner = 'SCOTT';
```

OWNER	OBJECT_NAME	TRACK_TIME	DML	SCAN	IDX
-----	-----	-----	-----	-----	-----
SCOTT	TEST	13-MAY-2018	YES	YES	NO
SCOTT	TEST	14-MAY-2018	NO	YES	NO

Notes

- Accuracy of a day.
- Index segments are also tracked
- Objects in 'SYSTEM', 'SYSAUX' tablespaces are not tracked
- Scans are tracked in UGA and periodically flushed to SGA.
- Can be turned off at session level

Row level heat map

- Track row modification. Rolled up to block level
- Package `dbms_heat_map` has API's to provide block level heat map information
- Used by ADO to filter 'hot' blocks

Example

- ```
select tablespace_name, relative_fno, block_id,
 writetime from
table(dbms_heat_map.block_heat_map('SCOTT','TEST'))
;
```

| TABLESPACE | RELATIVE_FNO | BLOCK_ID | WRITETIME |
|------------|--------------|----------|-----------|
| TBS_1      | 5            | 265236   | 13-MAY-18 |

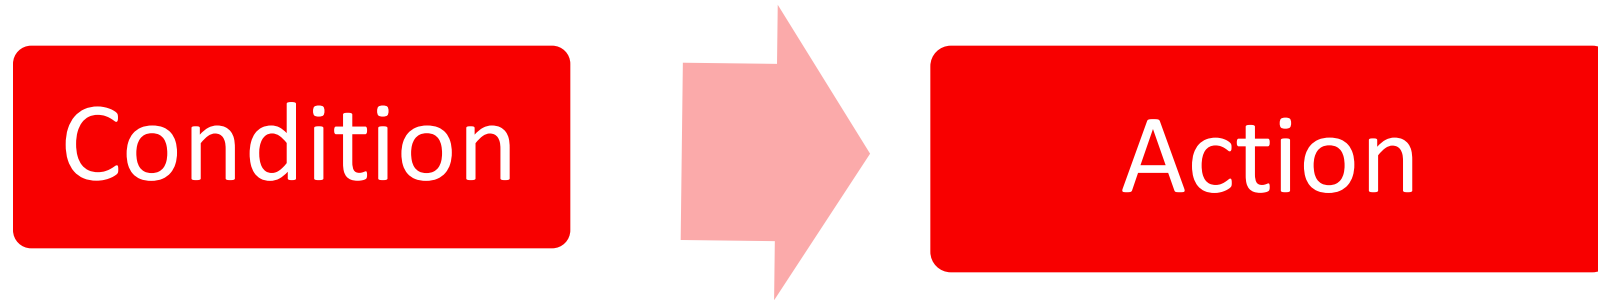
# Automatic Data Optimization (ADO)

- Automate the compression and movement of data during different stages of its lifecycle within the database
- Specify ILM rules using 'policies'
- Ability to create policies on tablespaces, tables, table partitions, table subpartitions
- Evaluate and execute policies automatically

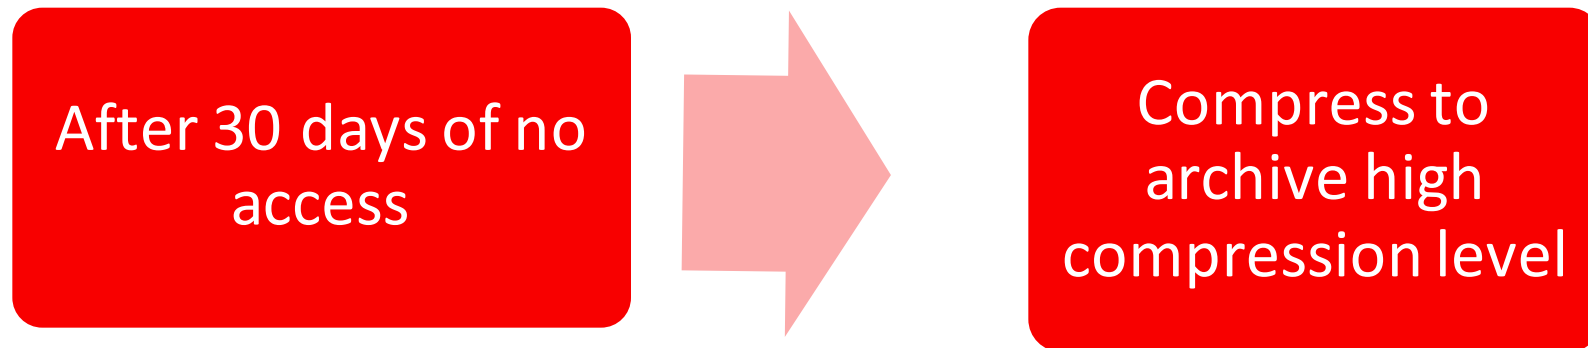
# ADO Interfaces

- SQL
  - E.g. Add policies, list policies on objects
- PL-SQL
  - E.g. Execute policies
- EM (not covered in this talk)
  - Enterprise Manager (EM) support to visualize heat-map, add policies, list policies etc.

## ADO policy



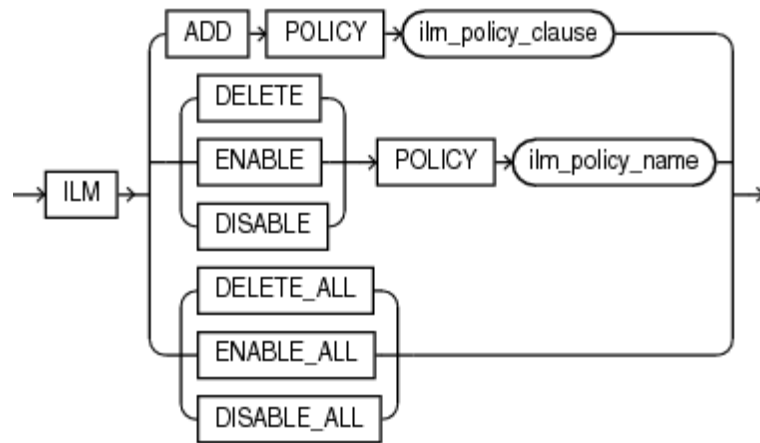
## Example



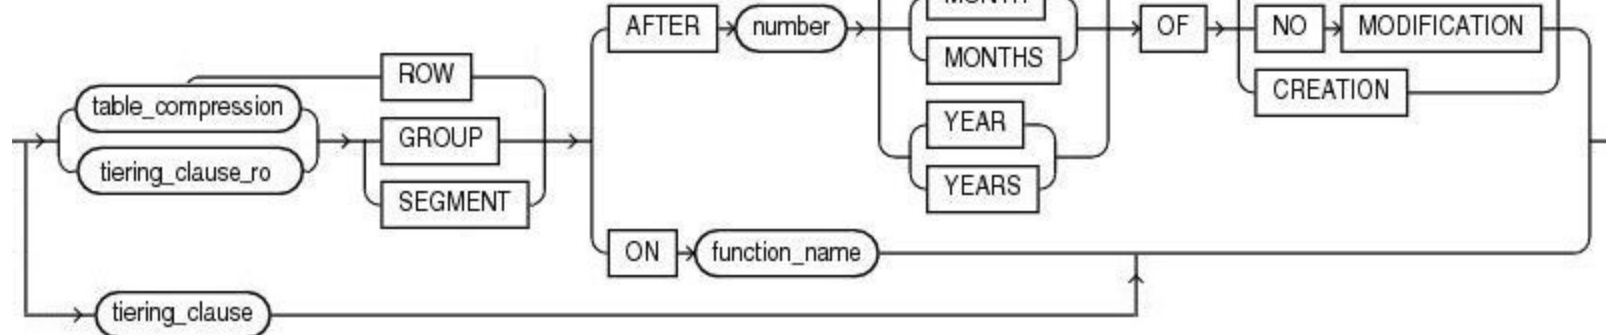
# ADO Syntax (SQL)

- An ADD policy SQL:
  - CREATE TABLE t1 (n int) ILM ADD POLICY *COLUMN STORE COMPRESS FOR QUERY HIGH AFTER 7 DAYS OF NO MODIFICATION*
- Read as:
  - Create a table and **add** an **ILM policy** to it such that the table **segment** is compressed to level **COLUMN STORE COMPRESS FOR QUERY HIGH** (HCC compression) after the segment has **not** been **modified** for **7 days**.

# ADO Syntax Diagram



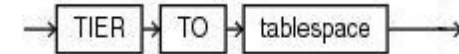
**Fig. A (ilm\_clause)**



**Fig. B (ilm\_policy\_clause)**

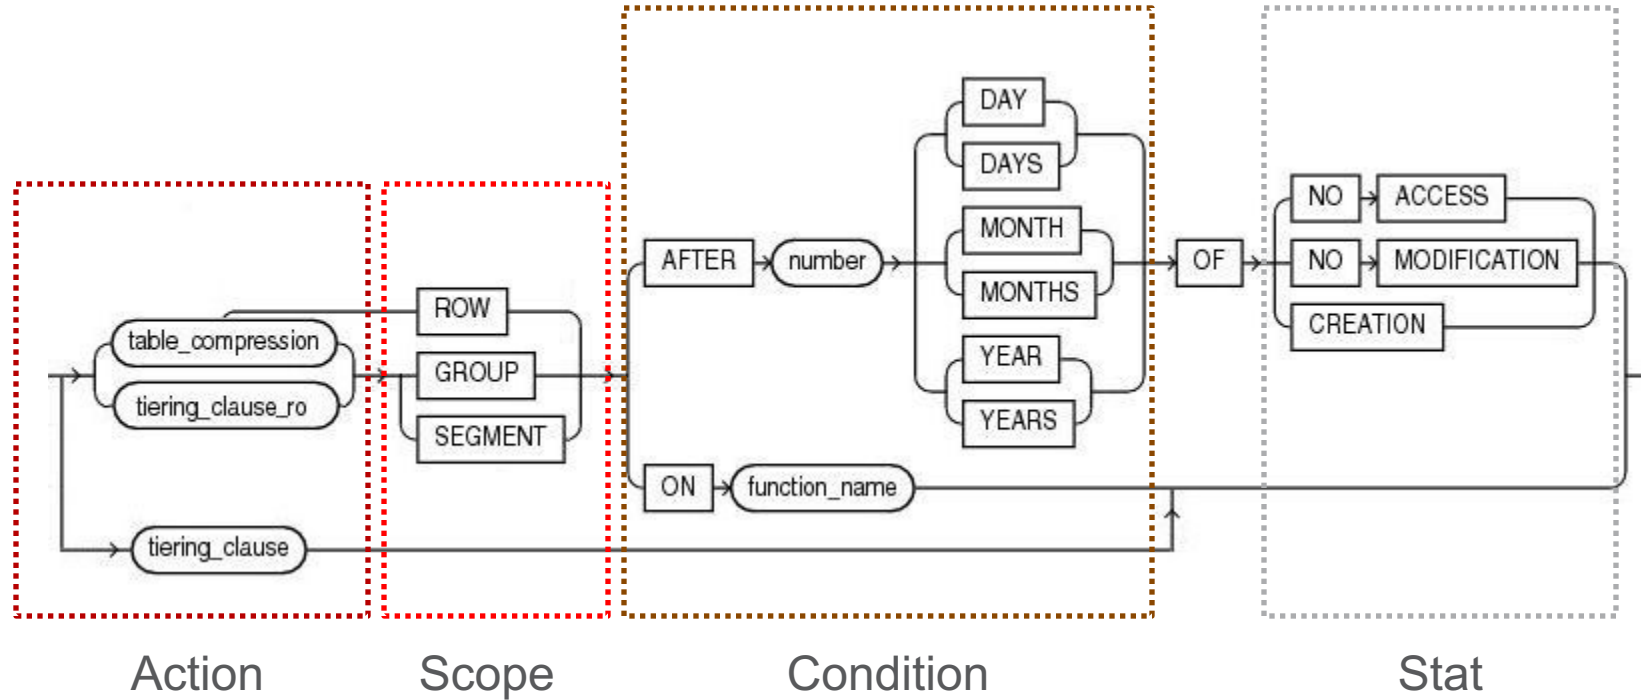


**Fig. C (tiering\_clause\_ro)**



**Fig. D (tiering\_clause)**

# ADO Syntax (Deconstruction)



**Fig. B (ilm\_policy\_clause)**

# ADO Syntax

- Policies can be specified on (via CREATE/ALTER):
  - Tablespace
    - `alter tablespace tbs_1 default ilm add policy column store compress for query high segment after 7 days of no modification;`
  - Table
    - `alter table emp ilm add policy row store compress advanced row after 3 days of no modification;`
  - Partition
    - `create table t1 (C1 number, C2 varchar2(9)) partition by list(C2) (partition p1 values('clerk', 'salesman') ilm add policy column store compress for archive high segment after 3 months of creation);`
  - Subpartition

# ADO Policy Details Views

- (USER/DBA)\_ILMPOLICIES – ADO policies and their status
- (USER/DBA)\_ILMDATAMOVEMENTPOLICIES – Details of all ADO policies
- (USER/DBA)\_ILMOBJECTS – Policy associations with objects, inheritance information, and the status

# ADO Policy Comments

- Policy name is auto-generated
- Policies are additive - An object can have several policies
- Conflict resolution rules for conflicting policies
- Policies can be inherited by objects.
  - Inheritance rules generally follow Oracle's existing model of inheritance (e.g. compression)

# Compression Policy

- Several levels of compression:

Advanced row compression

```
ROW STORE COMPRESS
ADVANCED
```

Warehouse compression (Hybrid  
Columnar Compression)

```
COLUMN STORE COMPRESS
FOR QUERY [LOW|HIGH]
```

Archive compression (Hybrid Columnar  
Compression)

```
COLUMN STORE COMPRESS
FOR ARCHIVE [LOW|HIGH]
```

# Compression Policy

- Implicit hierarchy among the compression levels
- An object can have a compression policy at each compression level provided
  - The policy condition honors the hierarchy of compression levels
  - All policies are on the same stat
- E.g.

**alter table t1 ilm add policy column store compress for query high segment after 10 days of no modification**

**alter table t1 ilm add policy column store compress for archive high segment after 5 days of no modification (Not allowed)**

# Incremental compression

- Benefits workloads through delayed background compression
- Uses heat map to compress only 'cold' blocks
- E.g.

```
alter table t1 ilm add policy row store compress advanced
row after 10 days of no modification
```

# Storage Tiering Policy

- An object can have only one storage tiering policy
- Syntax:
  - No condition clause
- Condition controlled via two system-wide ADO parameters, TBS\_PERCENT\_USED and TBS\_PERCENT\_FREE

```
create table t1 (c1 int) tablespace tbsHighSpeed ilm add
policy tier to tbsLowCost;
```

- Read as: Tier table t1 to tbsLowCost when tbsHighSpeed's usage goes beyond TBS\_PERCENT\_USED
- Move segments to destination tablespaces until freeness of the source tablespace (tbsHighSpeed) hits TBS\_PERCENT\_FREE

# GROUP keyword

- Way to specify policy action on dependent objects in addition to the object itself
- For policies with GROUP keyword, database would compress/move associated indexes and lob segments as well
- Pre-defined mapping from table compression levels to index compression level and LOB compression level

# Custom Policy

- Policy specification based on custom conditions
- Custom conditions encapsulated in a PL/SQL function
- Execution happens based on the truth value of the specified PL/SQL function
- Syntax:
  - `CREATE OR REPLACE FUNCTION business_logic (objn IN NUMBER) RETURN BOOLEAN;`
  - `ALTER TABLE t1 ILM ADD POLICY COLUMN STORE COMPRESS FOR QUERY LOW SEGMENT ON business_logic;`
- Can not be specified on tablespace or at ROW scope

# Policy Evaluation

- Policy evaluation and execution takes place in the system's maintenance windows
- An MMON task evaluates the policies on objects via MMON slaves by consulting the heat-map
- Results of evaluation are recorded in the view (USER/DBA)\_ILMEVALUATION\_DETAILS

# Policy Execution

- Objects whose policy conditions are satisfied, qualify for execution
- Execution engine creates jobs for qualifying policies.
- Jobs are run using `DBMS_SCHEDULER`
- Internally the jobs use `DBMS_REDEFINITION` and `alter table move partition DDL` for segment level operations and internal drivers for incremental compression

# Evaluation/Execution Views

- (USER/DBA)\_ILMTASKS – Tasks and their status. A task ID tracks an ADO evaluation/execution instance
- (USER/DBA)\_ILMEVALUATIONDETAILS – Evaluation details for each task .
- (USER/DBA)\_ILMRESULTS – Status and results of every ADO job.

# DBMS\_ILM\_ADMIN

- **Customize settings for ADO**

- `DBMS_ILM_ADMIN.CUSTOMIZE_ILM(DBMS_ILM_ADMIN.TBS_PERCENT_USED, 80);`
  - `DBMS_ILM_ADMIN.CUSTOMIZE_ILM(DBMS_ILM_ADMIN.TBS_PERCENT_FREE, 30);`

- **Disable/enable ADO**

- `dbms_ilm_admin.disable_ilm;`
  - `dbms_ilm_admin.enable_ilm;`

# DBMS\_ILM

- Immediate evaluation/execution of ADO policies without waiting for maintenance windows(Testing)

```
declare

v_executionid number;

begin

dbms_ilm.execute_ILM (ILM_SCOPE => dbms_ilm.SCOPE_SCHEMA,
 execution_mode => dbms_ilm.ilm_execution_online,
 task_id => v_executionid);

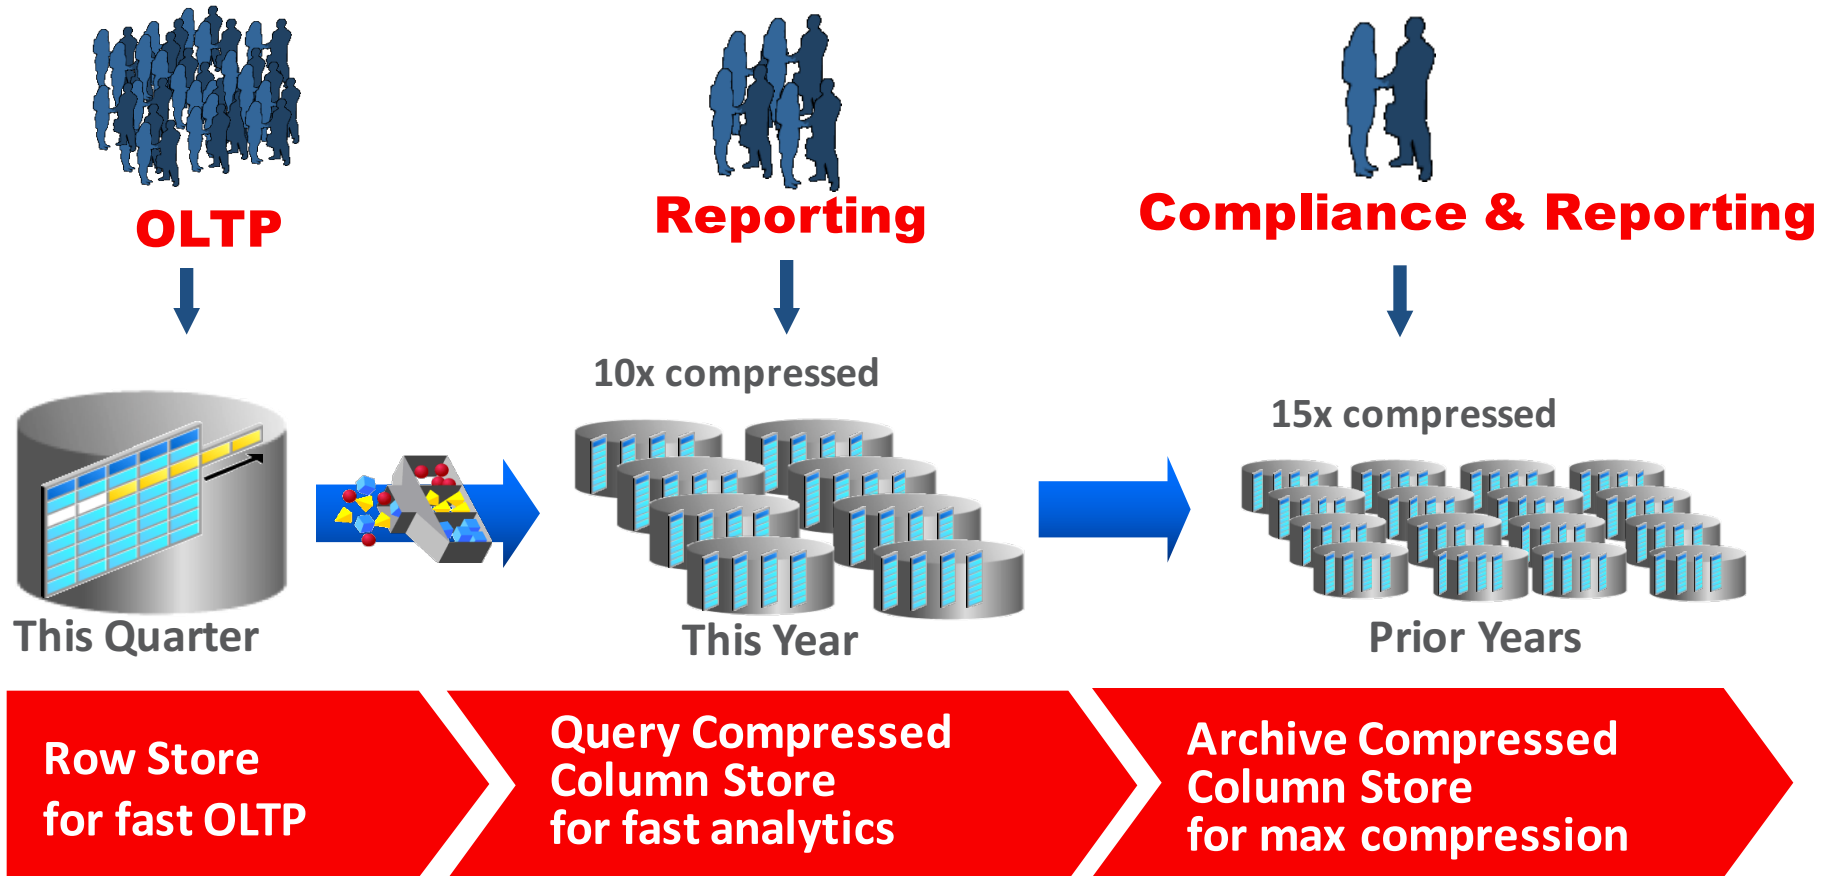
end;
/
```

# ADO implementation pointers

- Create ADO policies to reflect ILM strategy
  - Use heat map to monitor and understand usage patterns
  - Create candidate policies
  - Use the `dbms_ilm` package to evaluate policies, preview ADO actions and execute the candidate policies
  - Use `dbms_ilm_admin` package to control ADO execution environment

# Automatic Data Optimization

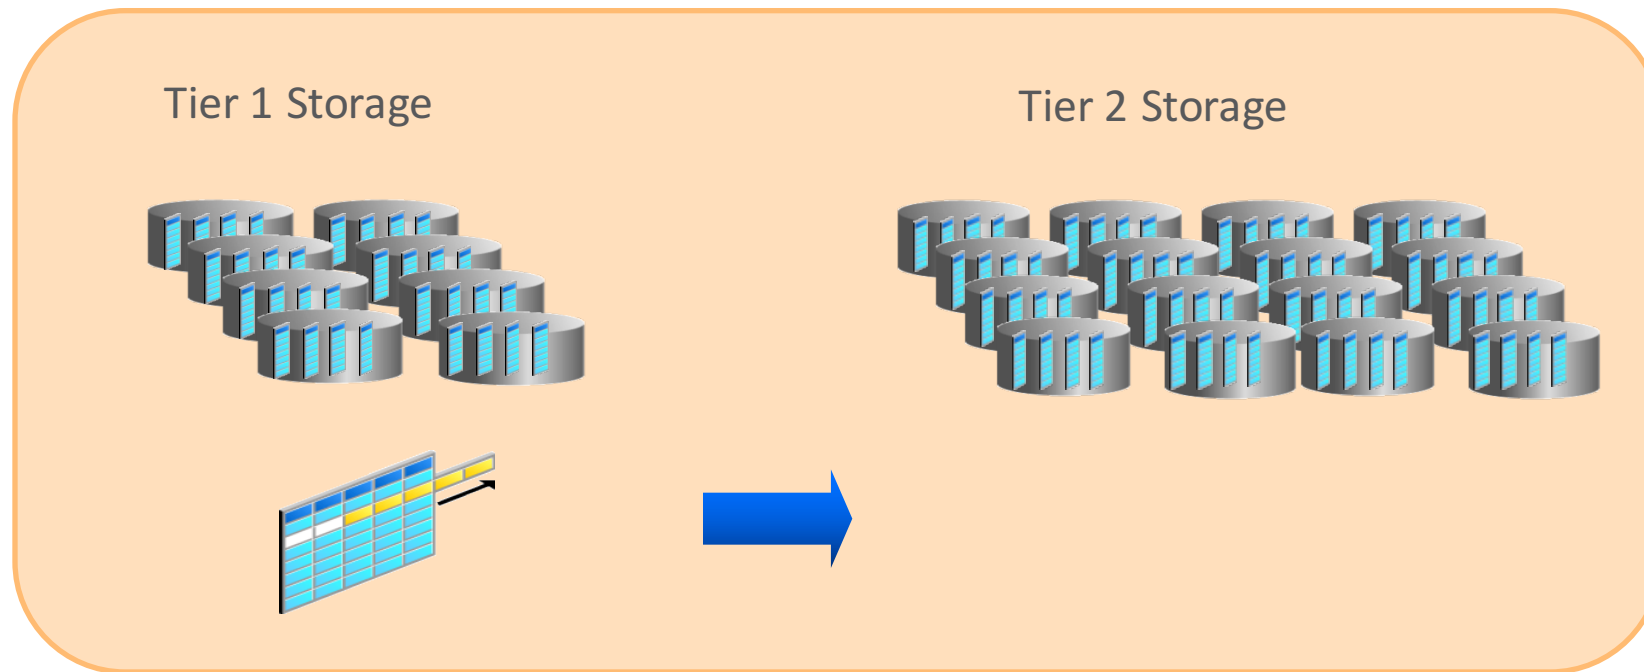
## Best Practice Example



As data cools down, Automatic Data Optimization can automatically convert Advanced Row compressed data to Columnar compressed online

# Automatic Data Optimization

## Usage Based Storage Tiering



**As storage pressure increases in Tier 1 storage, segments with tiering policies defined will automatically move to Tier 2 storage. If partitioned, there will be no application outage**

ORACLE®