

ORACLE®

Oracle 10g Performance Enhancements / Features on zLinux

Thomas Niewel
Principal Sales Consultant
Business Unit Database
Oracle Germany

Oracle10g Manageability Features

• Infrastructure

- Easy OracleJVM Monitoring
- SYSAUX Tablespace
- Enhanced Wait Model
- Database Features Usage Tracking
- Enhanced Database Time Model
- Top Consumers Performance Monitoring & Diagnostic
- JDBC: End-to-end Java/J2EE Tracing
- SQLJ: End-to-end Java/J2EE Tracing
- Automated Routine Administration Tasks
- Proactive Server-based Alerts and Notifications
- Database Statistics and Workload Respository
- Threshold-Based Alerts

• Server Configuration

- Automated Storage Management Configuration
- Automated RAC Services Configuration
- Simplified Upgrade for RAC and OPS Databases
- Automated Enterprise Manager Configuration
- Automated Portable Clusterware Installation
- Automated Configuration of Recovery Area
- Out-of-the-box LDAP Configuration
- Simplified Initialization Parameters
- Easy Upgrade
- Upgrade Information Tool
- Simplified Database Install
- User Default Tablespace

• Instance Tuning

- User-Initiated Buffer Cache Flushing
- Database Resource Manager - Adaptive Consumer Group Mapping
- Database Resource Manager - Fixed CPU Quota
- New Performance Overview Charts in Oracle Enterprise Manager
- Improved SQL Reporting Using Oracle Enterprise Manager
- Integration of Resource Manager and Profiles
- Automated Checkpoint Tuning
- Transaction Rollback & Recovery Monitoring
- Database Performance Analyzer
- Redo Logfile Sizing Advisor
- Automated SGA Memory Tuning

• Application Tuning

- SQLAccess Advisor
- Materialized View Tuning API
- SQL Tuning Advisor
- Enhanced SQLAccess Advisor
- Automatic Optimizer Statistics Collection

• Backup and Recovery

- Enhanced RMAN Reporting
- Backup Compression
- Bounded Backup Window
- Manage RMAN Scripts
- Recovery Area Full Alert

• Storage Management

- Multiple Default Temporary Tablespace Support for SQL Operations
- Automated Storage Management
- Rename Tablespace

Oracle 10g on z/Linux

- The focus of this presentation is z/Linux (and a little bit generic functionality)
 - Sizing the Database Cache
 - Async I/O
 - XSTOR Page Migrations if running under VM
 - Some important Oracle 10g features

Oracle 10g on z/Linux Storage issues

Oracle 10g on z/Linux

Oracle Enterprise Manager (SYSTEM) - Operating System Details - Microsoft Internet Explorer

Address: <http://192.168.7.32:5500/em/console/ecn/track/hc/view/osGeneral?target=lxora.boeblingen.de.ibm.com&type=host?event=doLoad>

ORACLE Enterprise Manager 10g
Database Control

Host: lxora.boeblingen.de.ibm.com > Operating System Details
Data Collected Mar 13, 2005 7:50:48 AM

Operating System Details

Host **lxora.boeblingen.de.ibm.com**
Operating System **SuSE SLES-8 (S/390X) 2.4.21 261 (64-bit)**
Vendor **SuSE**

General | [File Systems](#) | [Packages](#)

General Information

Distributor Version **N/A**
Maximum Swap Space (MB) **2047.984**

Operating System Properties

Previous 1-25 of 429 Next 25 History

Name	Source	Value
Maximum Swap Space (MB)	/sbin/sysctl	2047.984
abi.fake_utname	/sbin/sysctl	0
abi.trace	/sbin/sysctl	0
apldata.interval	/sbin/sysctl	60
apldata.timer	/sbin/sysctl	0
dev.raid_speed_limit_max	/sbin/sysctl	100000
dev.raid_speed_limit_min	/sbin/sysctl	100
fs.aio-max-pinned	/sbin/sysctl	38400
fs.aio-max-size	/sbin/sysctl	131072
fs.aio-pinned	/sbin/sysctl	300
fs.dir-notify-enable	/sbin/sysctl	1
fs.file-max	/sbin/sysctl	65536
fs.lease-break-time	/sbin/sysctl	45
fs.leases-enable	/sbin/sysctl	1

Oracle 10g on z/Linux

Sizing the Database Cache – 64 Bit Architecture

- Less physical I/O
- Helps to improve the performance
 - I/O is 2 ms (I/O Subsystem cache)
 - 5-15 ms (physical I/O)
 - # of I/O's per disk is limited

Oracle 10g on z/Linux

Async I/O

- Was not supported with Oracle 9.2 on z/Linux
- Improves I/O Performance
- Oracle has to be relinked to support Async I/O

```
zLinux1:~>cd $ORACLE_HOME/rdbms/lib
```

```
zLinux1:~/oracle010g/rdbms/lib> make -f ins_rdbms.mk async_on
```

```
zLinux1:~/oracle/O10g/rdbms/lib> make -f ins_rdbms.mk ioracle
```


Oracle 10g on z/Linux

Async I/O changes on init.ora.

- Raw devices

- DISK_ASYNC_IO = TRUE

- Filesystem

- FILESYSTEMIO_OPTIONS=ASYNCH

- FILESYSTEMIO_OPTIONS=SETALL

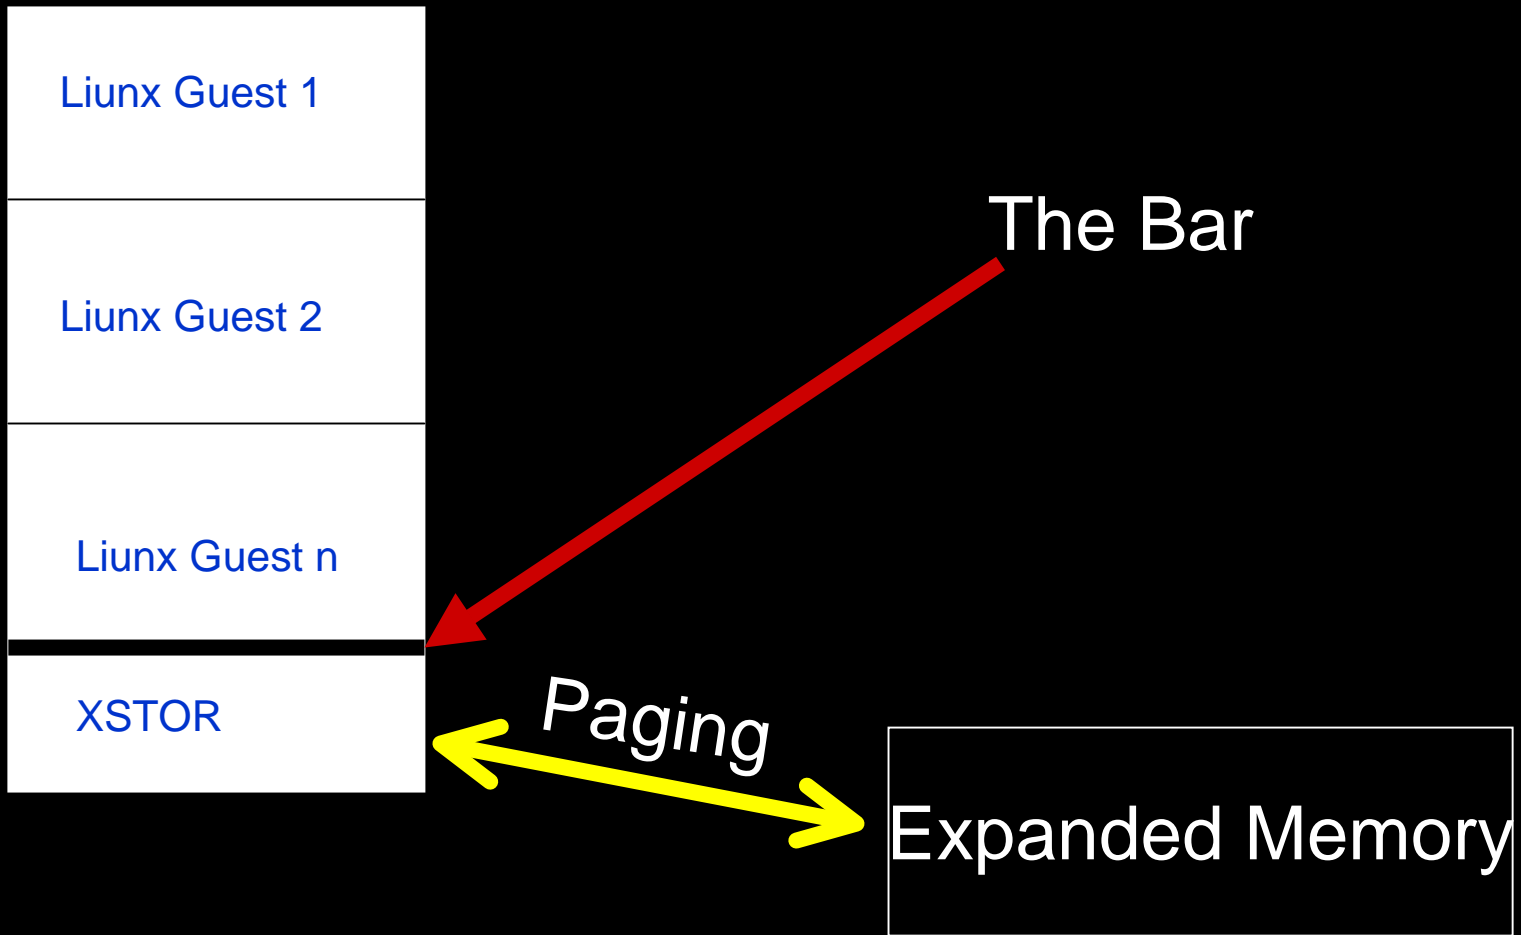
- Async I/O + Direct I/O

Oracle 10g on z/Linux

- Async-I/O
 - Process passes I/O-request to OS
 - work is processed in parallel
 - OS sends an interrupt when the work is finished
- Direct-I/O
 - Unix Filesystem Buffer Cache is not used
 - no double buffering

Oracle 9.2 / 10g on z/Linux

z/VM Storage Management



Oracle 9.2 / 10g

How can these problems be solved ?

- wait for a newer VM Release (5.2 ?)
- Use LPAR's
- Oracle Parameters (10g under investigation)
 - LOCK_SGA=TRUE
 - root authorization required
 - gpasswd -a root dba
 - chown root \$ORACLE_HOME/bin/oracle
 - PRE_PAGE_SGA=TRUE

Some generic stuff

Automatic Statistics Gathering

Reasons for automatic statistic gathering

- It is not an easy task to generate “correct object statistics”
 - Analyze time
 - There are objects which have to be analyzed rarely
 - “Sample Size” too big
 - Statistics are inaccurate
 - “Sample Size” too small
- It’s not easy to determine columns for histograms

“Statistic Gathering” Job in 10g

- Statistics are **automatically generated**
- Predefined DB Scheduler Job
 - Scheduler Program and Job-definitions are defined in catmwin.sql (Called from catproc.sql)
 - Runs in a predefined window (MAINTENANCE_WINDOW_GROUP)
Default: 10pm-6am (Mo-Fr), 12am-12am (Sa)
 - Scheduler terminates Job, if it doesn't end in that window
endet (STOP_ON_WINDOW_CLOSE)
 - Restartable
- Manual Statistic gathering is only necessary in case of :
 - Bulk loads (e.g. in DW)
 - “Volatile” Tables

View: DBA_SCHEDULER_JOBS

```
SQL> SELECT job_name, state, comments FROM  
dba_scheduler_jobs;
```

JOB_NAME	STATE

COMMENTS	

PURGE_LOG	SCHEDULED
purge log job	
GATHER_STATS_JOB	SCHEDULED
system default gather statistics maintenance job	

Statistic Gathering Job

- PL/SQL Procedure:
DBMS_STATS.GATHER_DATABASE_STATS_JOB_PROC
- Gather statistics for
 - Objects with missing or STALE Statistics
 - Dictionary and User Objects
- Determines
 - Objects with STALE statistics
 - Sample size
 - Histograms
- Does not collect statistics for
 - Global temporary Objects
 - Materialized View Logs
 - Objects with lock_table_stats

Objects with STALE Statistics

- Modification of more than 10 pct of rows in table or the table was truncated since the last Analyze
- Modification Monitoring:
 - # of row modifications () since the last Analyze
 - Operations
 - Insert
 - Delete
 - Update
 - Truncate

Monitoring: DML-Usage

- Determined by DML-Driver during execution
- Overhead : less 1%
- Prerequisite : STATISTICS_LEVEL =TYPICAL
- Catalog view : DBA_TAB_MODIFICATIONS

Monitoring DML Usage

```
SQL> SELECT * FROM user_tab_modifications;
```

TABLE_NAME	PARTITION_NAME	SUBPARTITION_NAME	INSERTS	UPDATES	DELETES
PLSQL_PROFILER_RUNS			2	4	0
04.11.04	NO	0			
PLSQL_PROFILER_UNITS			8	119	0
04.11.04	NO	0			

Column Usage Monitor

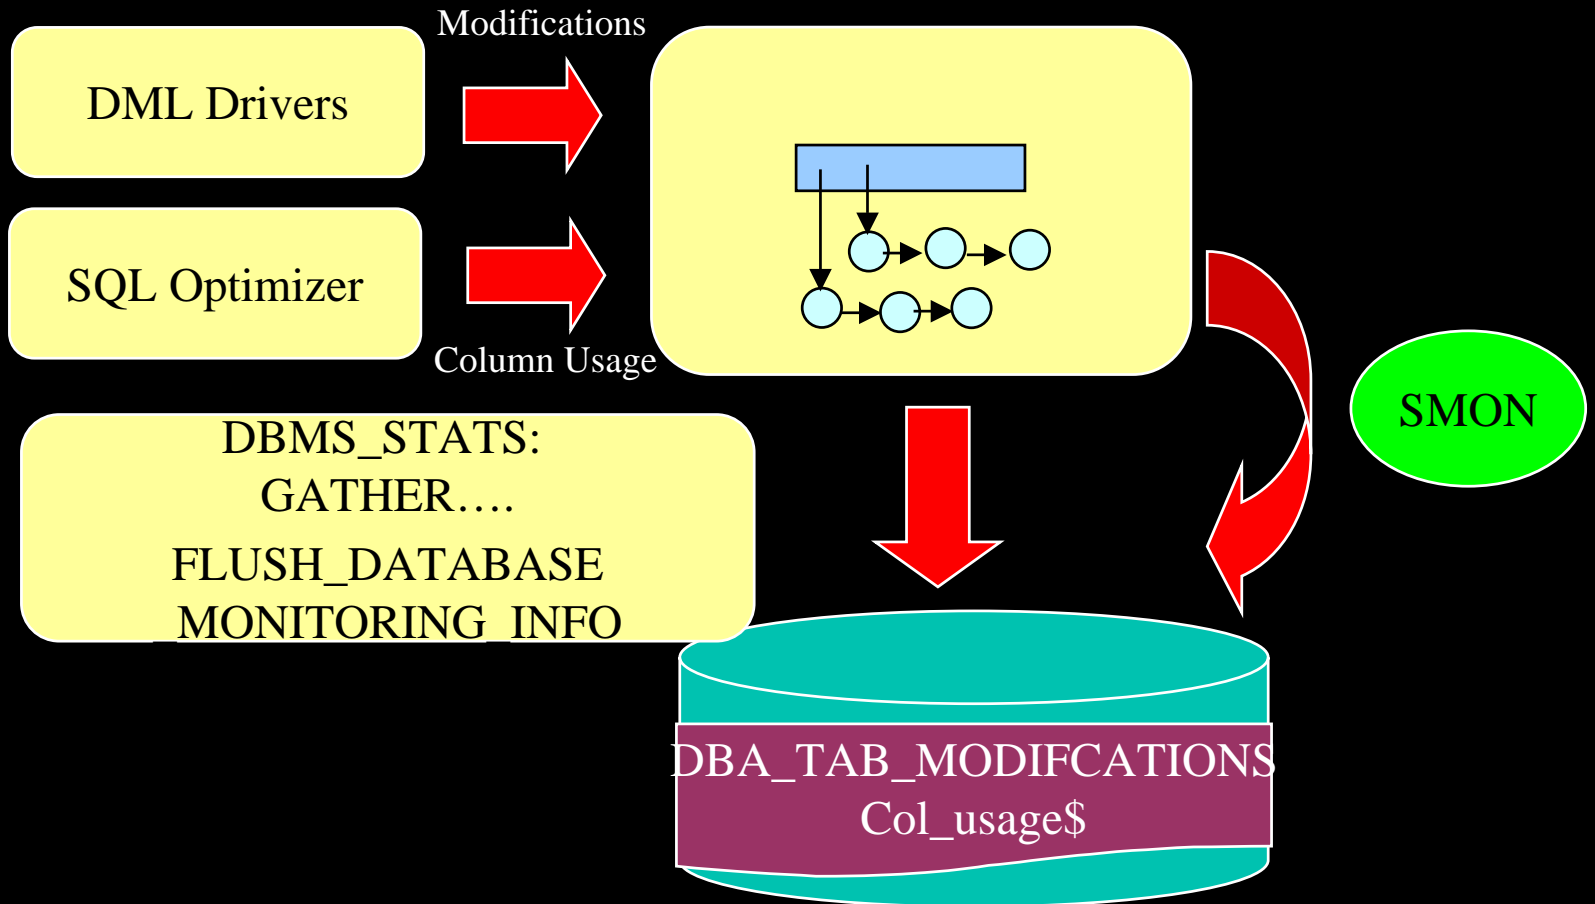
- Saves Information of Column usage in
 - Equality
 - Equi-join
 - Range
 - Non Equi-join
 - Like
 - IS (NOT) NULL

Column usage Monitor

- Is saved by Query Optimizer during compile
- Information is stored in col_usage\$

```
OBJECT_NAME          INTCOL#  TIMESTAMP  EQUALITY_PREDS
-----
NONEQUIJOIN_PREDS  RANGE_PREDS  LIKE_PREDS  NULL_PREDS
-----
PLSQL_PROFILER_DATA 3          04.11.2004 3
0                      0          0          0
```

Monitoring Mechanism



Tuning in Oracle 10g

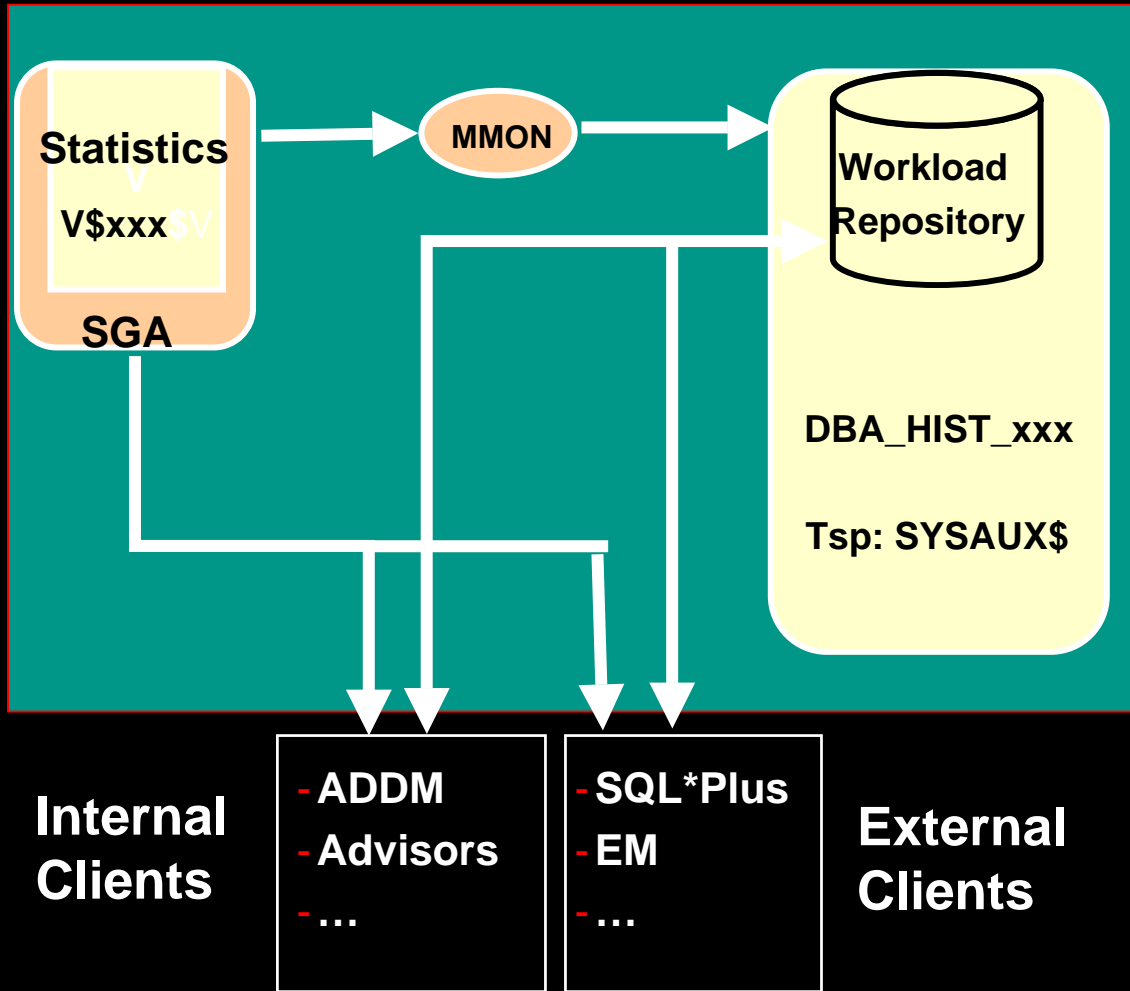
Traditional Tuning

- Maintain Object Statistics
- Create/drop indexes
- Statspack & utlbstat / utlestat
- Optimization of Init.ora parameters
- Analyze of Explain plan's and TKPROF reports
- Optimize SQL

New Tuning mechanisms

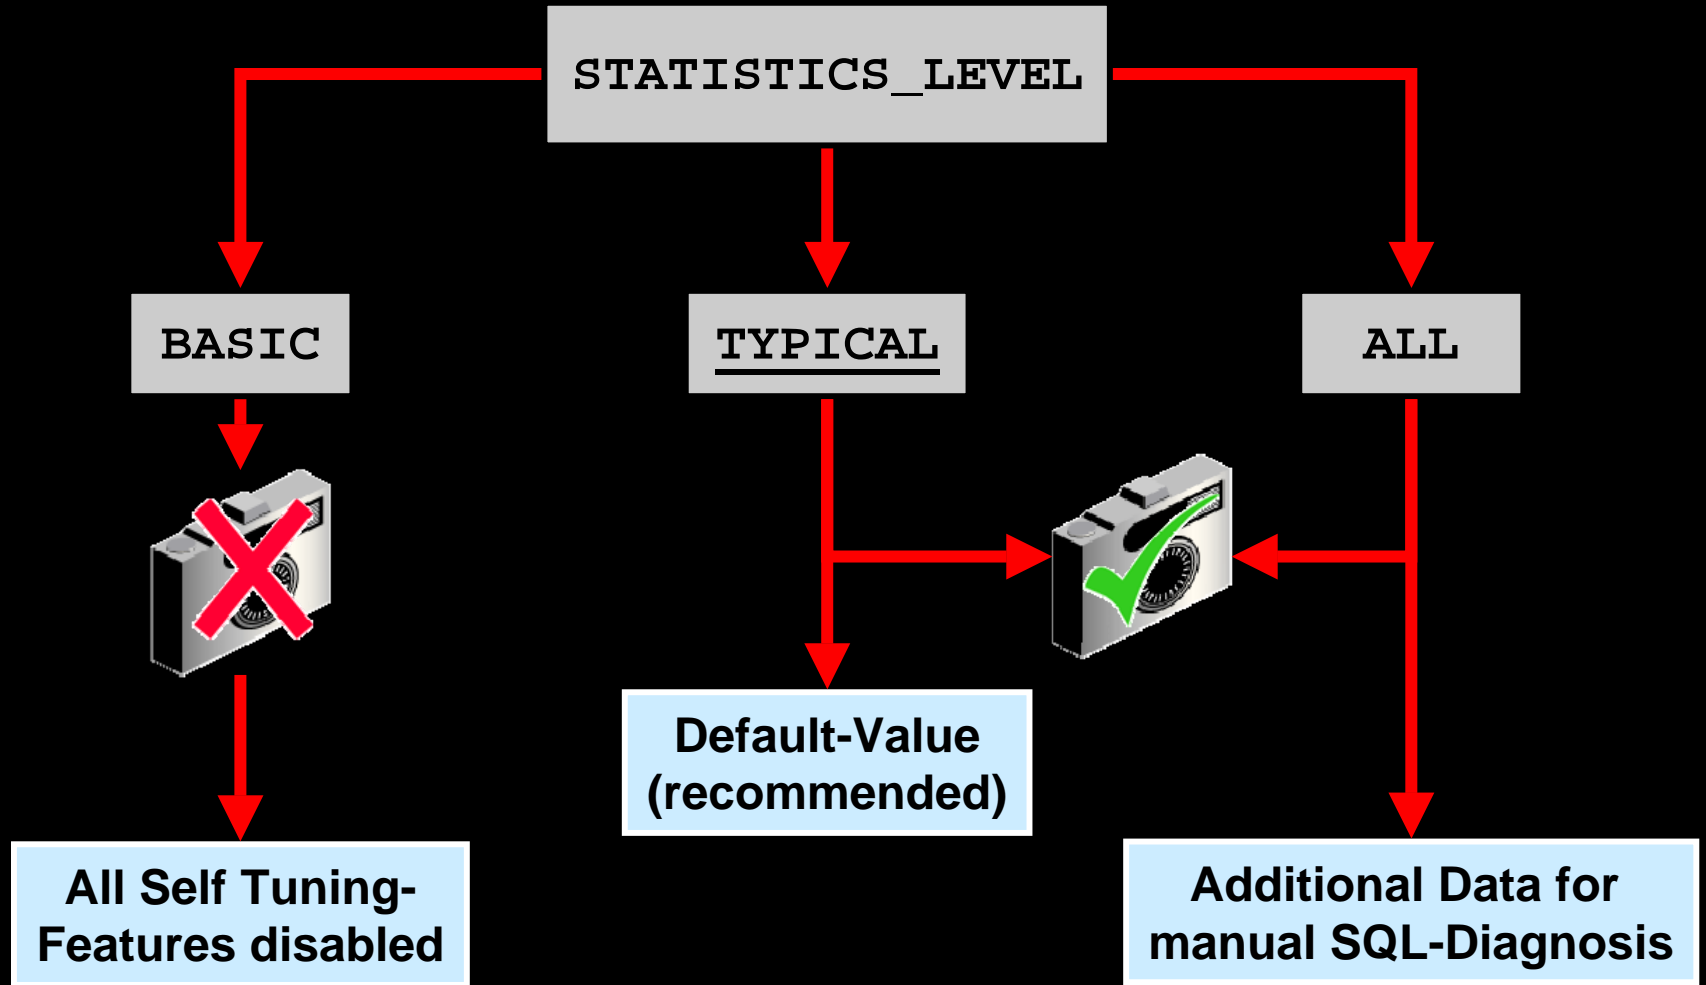
- ADDM
- Automatic Advisors
 - SQL Tuning Advisor
 - SQL Access Advisor
- AWR - is basis for all Advisors

Automatic Workload Repository (AWR)



- Base Statistics, Metrics, SQL-Statistics, Active Session History
- Automatic Snapshots (Default 1h)
- “Historic” Data (Default 7 days)
- Automatic Space Management
- “Light Weight-Capture”

AWR: Configuration



AWR

- Collects Information
 - Wait events occurred
 - Time model statistics indicating the amount of DB time associated with a process from the V\$SESS_TIME_MODEL and V\$SYS_TIME_MODEL views.
 - Active Session History (ASH) statistics
Source(V\$ACTIVE_SESSION_HISTORY)
 - System/session statistics from the V\$SYSSTAT and V\$SESSTAT views.
 - Object usage statistics.
 - Resource intensive SQL statements.

AWR

- Per default snapshots are taken every hour and retained for 7 days.
 - Interval can be altered

```
BEGIN DBMS_WORKLOAD_REPOSITORY.modify_snapshot_settings
  (retention => 43200, -- Minutes (= 30 Days).
   interval => 30);    -- Minutes. Current value retained if NULL.
END;
```

- Snapshots can be taken manually

```
EXEC DBMS_WORKLOAD_REPOSITORY.create_snapshot;
```

AWR

- The pair of snapshots associated with a baseline are retained until the baseline is explicitly deleted:

```
BEGIN
  DBMS_WORKLOAD_REPOSITORY.create_baseline (
    start_snap_id => 210, end_snap_id => 220,
    baseline_name => 'batch baseline'); END;
```

```
BEGIN
  DBMS_WORKLOAD_REPOSITORY.drop_baseline (
    baseline_name => 'batch baseline', cascade
    => FALSE);
END; /
```


AWR – dictionary views

- V\$ACTIVE_SESSION_HISTORY - Displays the active session history (ASH) sampled every second.
- V\$METRIC - Displays metric information.
- V\$METRICNAME - Displays the metrics associated with each metric group.
- V\$METRIC_HISTORY - Displays historical metrics.
- V\$METRICGROUP - Displays all metrics groups.
- DBA_HIST_ACTIVE_SESS_HISTORY - Displays the history contents of the active session history.
- DBA_HIST_BASELINE - Displays baseline information.
- DBA_HIST_DATABASE_INSTANCE - Displays database environment information.
- DBA_HIST_SNAPSHOT - Displays snapshot information.
- DBA_HIST_SQL_PLAN - Displays SQL execution plans.
- DBA_HIST_WR_CONTROL - Displays AWR settings.

AWR – Reports

- `@$ORACLE_HOME/rdbms/admin/awrrpt.sql`
- `@$ORACLE_HOME/rdbms/admin/awrrpti.sql`
- OEM

AWR – Reports

Top 5 Timed Events

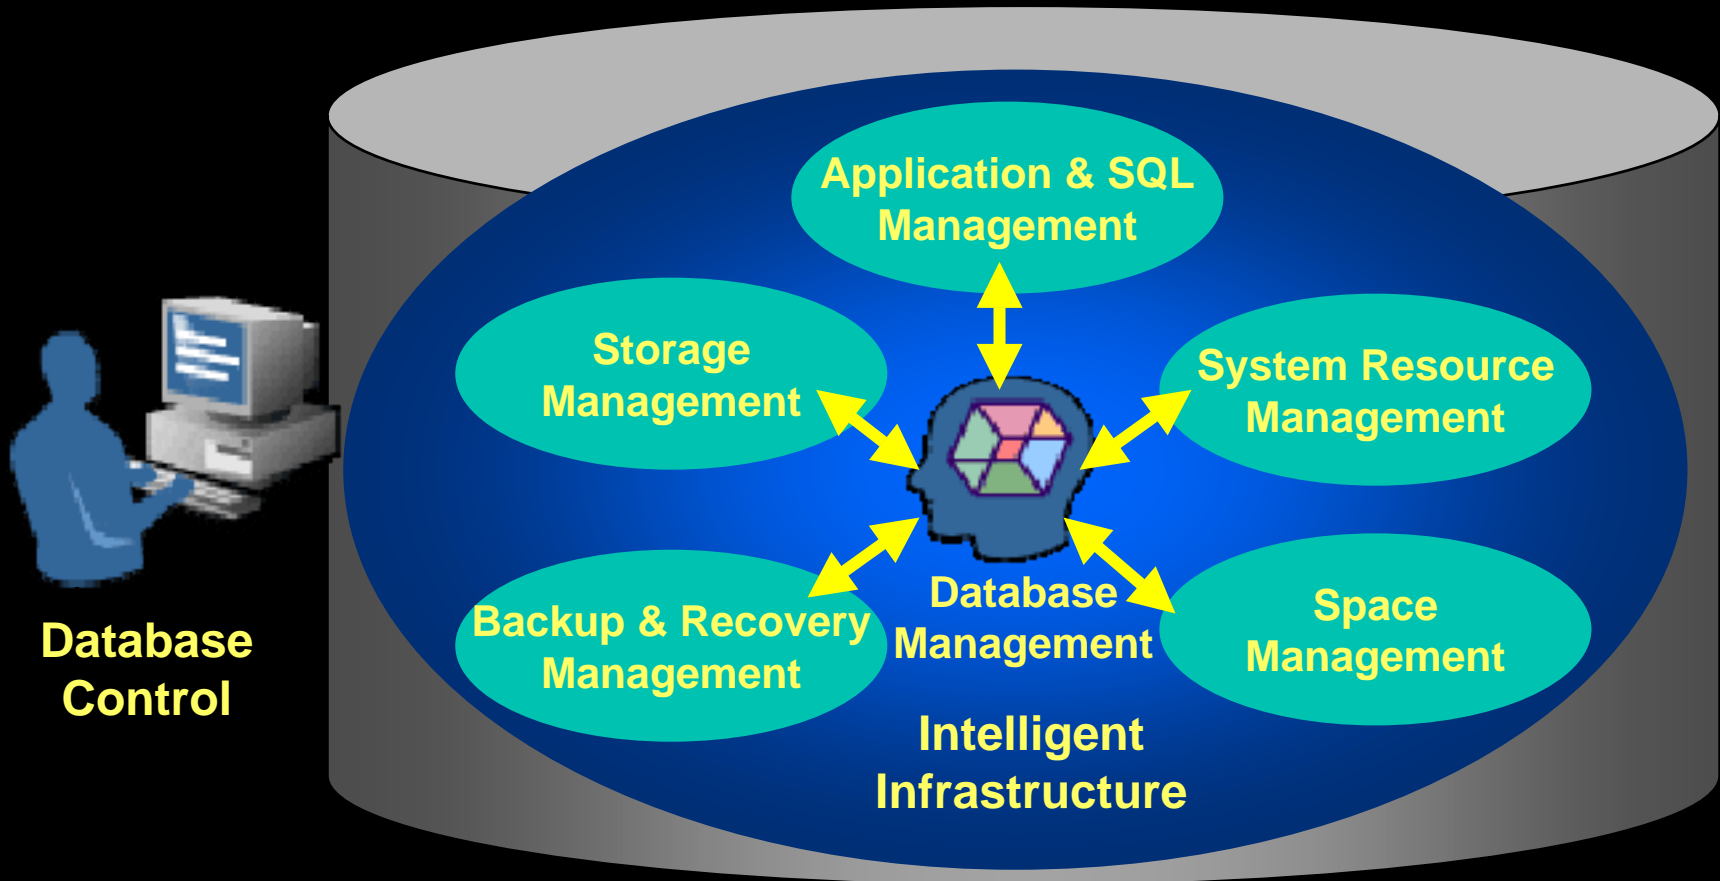
Event	Waits	Time(s)	Percent Total DB Time	Wait Class
db file sequential read	2,820	19	62.34	User I/O
control file parallel write	3,623	15	47.96	System I/O
CPU time		12	38.05	
class slave wait	1	5	16.25	Other
db file scattered read	364	3	10.40	User I/O

Main Report

- [Report Summary](#)
- [Wait Events Statistics](#)
- [SQL Statistics](#)
- [Instance Activity Statistics](#)
- [IO Stats](#)
- [Buffer Pool Statistics](#)
- [Advisory Statistics](#)
- [Wait Statistics](#)
- [Undo Statistics](#)
- [Latch Statistics](#)
- [Segment Statistics](#)
- [Dictionary Cache Statistics](#)
- [Library Cache Statistics](#)
- [SGA Statistics](#)
- [Resource Limit Statistics](#)
- [init.ora Parameters](#)

[Back to Top](#)

Oracle Database 10g – Self-Managing Database



More Informationen

- Note 276103.1: PERFORMANCE TUNING USING 10g ADVISORS AND MANAGEABILITY FEATURES
- Note 271196.1: Automatic SQL Tuning - SQL Profiles
- Note 262687.1: How to use the SQL Tuning Advisor
- Oracle® Database Performance Tuning Guide

Q&A

ORACLE®