

DATABASES

ADVANCED

3 DAYS

Oracle Database Performance Tuning

Three days of evidence-led Oracle tuning for DBAs and developers who need slow databases explained rather than guessed at. Covers a repeatable tuning method, AWR, ADDM and ASH, wait events, the cost-based optimiser and execution plans, statistics, indexing, joins, SQL tuning and plan stability, memory and I/O tuning, and contention resolution. Delivered on Oracle 19c and 23ai.

Why Oracle Performance Tuning Training?

A method, not a checklist

Teams learn a repeatable diagnostic sequence that starts from where time is actually spent rather than from a list of parameters to change.

Plan stability included

SQL plan baselines and profiles are covered, so improvements survive the next statistics gathering run.

Wait events as the entry point

Every investigation begins with wait analysis, which is what separates a tuned database from a lucky one.

Licensing addressed honestly

The course is clear about which diagnostic features require the Diagnostics and Tuning Pack and offers alternatives where they are not licensed.

What you'll be able to do

- ✓ Apply a repeatable, evidence-led method to any performance problem
- ✓ Generate and interpret AWR, ADDM and Active Session History output
- ✓ Use wait event analysis to identify the dominant bottleneck
- ✓ Explain how the cost based optimiser produces and costs a plan
- ✓ Read execution plans including real time SQL monitoring output
- ✓ Diagnose cardinality misestimation and correct statistics problems
- ✓ Select access paths and design indexes for a given workload
- ✓ Recognise join methods and understand query transformations
- ✓ Stabilise plans using SQL plan baselines, profiles and hints
- ✓ Tune instance memory, I/O and parallel execution
- ✓ Diagnose and resolve latch, lock and buffer contention
- ✓ Identify application and design level causes that no parameter will fix

Who this is for

- Oracle database administrators responsible for performance
- Application and PL/SQL developers writing performance sensitive code
- Platform and infrastructure engineers supporting Oracle workloads
- Data engineers running large Oracle based batch processes

- Technical leads accountable for application response times

Curriculum

01 A Method for Tuning

- Defining the problem in measurable terms before touching anything
- Time based analysis: where is the database actually spending time
- Scope: instance wide, session specific or single statement
- Establishing a baseline and proving improvement afterwards
- Diagnostics and Tuning Pack licensing and unlicensed alternatives
- **Hands-on:** scope and baseline a supplied performance complaint

02 AWR, ADDM and Active Session History

- AWR snapshot mechanics, retention and report generation
- Reading an AWR report in the right order
- Top timed events, load profile and instance efficiency sections
- ADDM findings, recommendations and how much to trust them
- Active Session History for problems too short to appear in AWR
- Statspack as an alternative where packs are not licensed
- **Hands-on:** diagnose a workload from an AWR report alone

03 Wait Event Analysis

- Wait event classes and what each one points at
- CPU versus wait time and interpreting the balance
- I/O related events and separating storage from access path problems
- Concurrency, configuration and cluster wait classes
- Idle events and avoiding false leads
- **Hands-on:** trace a session and attribute its time to specific causes

04 The Optimizer and Execution Plans

- Cost based optimisation, cardinality and cost calculation
- Generating plans with EXPLAIN PLAN, DBMS_XPLAN and the cursor cache
- Estimated versus actual rows as the primary diagnostic signal
- Real time SQL Monitoring for long running statements
- Adaptive plans and adaptive statistics behaviour
- **Hands-on:** capture and interpret plans for a set of problem statements

05 Statistics and Cardinality

- Table, column, index and system statistics
- Histograms, skew and when they help or mislead
- Extended statistics for column groups and expressions
- Automatic statistics gathering, staleness and manual intervention
- Dynamic sampling and its cost
- **Hands-on:** create a misestimation, watch the plan collapse and correct it

06 Access Paths and Index Strategy

- Full table scan, index range scan, unique scan, skip scan and fast full scan
- B-tree, bitmap, function based and composite index design
- Clustering factor and why it decides index usefulness
- Index compression, invisible indexes and safe index testing
- Automatic Indexing behaviour and oversight
- **Hands-on:** redesign an index set and prove the plan and write cost changes

07 Joins and Query Transformation

- Nested loops, hash join and sort merge join and their conditions
- Join order selection and the optimiser's search space
- Subquery unnesting, view merging and predicate pushing
- Star transformation and its data warehouse applications
- Diagnosing a bad join choice from plan evidence
- **Hands-on:** correct a poorly joined multi-table query

08 SQL Tuning and Plan Stability

- SQL Tuning Advisor and SQL Access Advisor output
- SQL profiles and what they actually change
- SQL plan baselines for capturing and evolving known good plans
- Hints: appropriate use, common misuse and long term cost
- Bind variable peeking and adaptive cursor sharing
- **Hands-on:** stabilise a regressing statement using a plan baseline

09 Instance Tuning: Memory and I/O

- SGA and PGA sizing, automatic memory management and advisors
- Buffer cache behaviour, hit ratios and why they mislead
- Shared pool, library cache and hard parse reduction
- Redo, checkpoint and log writer tuning
- Storage latency targets, ASM layout and I/O calibration
- **Hands-on:** resize memory components using advisor evidence and measure the effect

10 Contention and Concurrency

- Enqueue waits, row lock contention and blocking sessions
- Buffer busy waits, hot blocks and ITL contention
- Latch and mutex contention and their usual causes
- Sequence, index leaf block and right-hand growth contention
- Deadlock diagnosis from trace files
- **Hands-on:** reproduce a contention hotspot and design it away

11 Application and Design Level Tuning

- Parsing, cursor management and connection pooling behaviour
- PL/SQL context switching and bulk processing gains
- Partitioning for pruning and parallel operations
- Parallel execution configuration and when it makes things worse
- Materialized views, query rewrite and result caching
- Building a tuning runbook and a regression watch process
- **Hands-on:** capstone tuning of an end to end workload with measured results