MODIFIABLE RFP TEMPLATE & VENDOR Q&A RUBRIC | VNA-RFP-050 | Industry Operating Platforms (Public Sector, B2B, POS & Transit)

Request for Proposals (RFP): Snowflake & Databricks Egress Sprawl to Oracle Autonomous Data Warehouse (ADW) & Select AI

Instructions for RFP Author: This template is structured so your IT, Architecture, and Procurement teams can issue a rigorous, apples-to-apples Request for Proposals. Customize the bracketed [Modifiable] fields below or use the companion .md file on the solution portal.

1. Issuing Organization & Procurement Schedule (Modifiable Fields)

RFP Field / Milestone	Modifiable Entry / Target Value
Issuing Organization & Sponsor	[Insert Legal Organization Name, Executive Sponsor & Primary Contact Email]
Project Scope Summary	Transition from High-Cost Snowflake / Databricks Compute Credits, Fragile CDC Copy Pipelines & Egress Fees to Oracle Autonomous Data Warehouse (ADW 23ai) + Zero-ETL Converged Lakehouse + Select AI + OAC
RFP Issue / Q&A; / Proposal Due Dates	[Issue: YYYY-MM-DD] | [Vendor Questions Due: +14 Days] | [Proposals Due: +30 Days]
Target Contract Type	Fixed-Price Milestone Deliverables with Automated Parity & Cutover Acceptance Gates

2. Mandatory Technical Scope of Work (Modifiable Requirements)

[REQ-ARCH-01] Automated Phase 0 AST & binary scan mapping High-Cost Snowflake / Databricks Compute Credits, Fragile CDC Copy Pipelines & Egress Fees artifacts directly to live Oracle dictionary metadata and profiling production data.

[REQ-ARCH-02] Target architecture built on Oracle Autonomous Data Warehouse (ADW 23ai) + Zero-ETL Converged Lakehouse + Select AI + OAC, consolidating cloned queries into governed core SQL views with runtime Smart Filters.

[REQ-ARCH-03] Automated generation of full Business Requirements Documents (BRD), Technical Approach Docs, Testing Specifications, and Verified Test Results.

[REQ-ARCH-04] Zero-trust SSO/MFA (Entra ID / Okta) paired with database-enforced row-level security (VPD / RAS) and SQLcl Liquibase CI/CD pipelines.

3. Component Conversion & Acceptance Matrix

Legacy Artifact / Workload	Required Target Architecture	Engineering Method & Acceptance Gate
Source Tier: High-Cost Snowflake / Databricks Compute Credits, Fragile CDC Copy Pipelines & Egress Fees	Oracle Autonomous Data Warehouse (ADW 23ai) + Zero-ETL Converged Lakehouse + Select AI + OAC	Automated AST & Metadata Uplift (85%-95%)
Primary Footprint: Monthly Cloud Data Warehouse & ETL Pipelines	Oracle APEX 24.2 Redwood UI + PL/SQL APIs	Binary-to-DB Metadata Mapping + Pilot Sprint
Workload & Queries: Analytical Tables, Marts & BI Dashboards	Canonical Core SQL Views + Runtime Smart Filters	Query Fan-Out Clustering (75%-90% Reduction)
Undocumented Legacy Business Rules & QA	Auto-Generated BRD, Approach Doc, Test Specs & Results	100% Automated Documentation & utPLSQL / Playwright

4. Mandatory Vendor Q&A; Questionnaire & Evaluation Benchmark Key

Bidders must answer each technical question below. Evaluation committee members can compare vendor answers directly against the Reference Benchmark Key:

Vendor Question 1: What are the primary levers that control the cost of Snowflake / Databricks to Oracle 23ai ADW? *[Score: ___ / 10]*

Reference Evaluation Benchmark: Cost is controlled by five levers: (1) collapsing cloned query fan-out across analytical tables, marts & bi dashboards into canonical core SQL views, (2) percentage of business logic already in database packages vs. UI triggers, (3) production data quality profiled during Phase 0, (4) external integration/SSO count, and (5) cutover downtime requirements.

Vendor Question 2: How does your tool convert binaries (.fmb, .rdf, etc.), map database metadata, and generate documentation? *[Score: ___ / 10]*

Reference Evaluation Benchmark: Our automated analyzer parses legacy binaries and source files, extracts every trigger, query, and UI control, cross-references them against live database dictionary metadata (ALL_TAB_COLUMNS, ALL_DEPENDENCIES), profiles actual column data distributions, and automatically generates the Business Requirements Document (BRD), Technical Approach Doc, Test Specifications, and Verified Test Results.

Vendor Question 3: What does a typical project timeline look like from kickoff to production cutover? *[Score: ___ / 10]*

Reference Evaluation Benchmark: Projects follow four gated phases: Phase 0 (Weeks 1-2: Binary AST Scan, Data Profiling & Auto-Generated BRD), Phase 1 (Weeks 3-6: Core Schema Uplift & Working Pilot), Phase 2 (Weeks 7-12: Automated Conversion Sprints & Test Spec Execution), and Phase 3 (Weeks 13-16: Parallel Run, Zero-Downtime Cutover & 30-Day Hypercare).

Vendor Question 4: How do you rationalize report/query fan-out and guarantee 100% customer code & data ownership? *[Score: ___ / 10]*

Reference Evaluation Benchmark: Vendor must demonstrate AST SQL clustering to collapse cloned queries into governed core views with runtime/AI-inference filters, and deliver 100% customer-owned SQL, PL/SQL, APEX, or cloud code with zero proprietary runtime lock-in.

5. Commercial Milestone Pricing & Evaluation Weighting (Modifiable)

Phase / Deliverable Gate	Required Acceptance Output	Bidder Fixed Fee	Eval Weight
Phase 0: AST Discovery & Blueprint	Complete code/schema AST inventory, dead-code manifest & core SQL fan-out map	$ [Insert]	20%
Phase 1: Core Schema & Logic Uplift	Target cloud/DB provisioning, automated DDL/code conversion & unit tests	$ [Insert]	30%
Phase 2: UI, Reports & UAT	Modernized UI/portals, runtime/AOP reports, security VPD/SSO & UAT sign-off	$ [Insert]	30%
Phase 3: Cutover & Hypercare	Zero-downtime cutover, row/hash reconciliation & 30-day production warranty	$ [Insert]	20%