VNA-SB-050 | Industry Operating Platforms (Public Sector, B2B, POS & Transit) | Snowflake / Databricks Egress Sprawl to Oracle 23ai ADW Consolidation

Snowflake & Databricks Egress Sprawl to Oracle Autonomous Data Warehouse (ADW) & Select AI

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 with automated binary/AST metadata profiling, deterministic cost-control levers, auto-generated BRD and test documentation, and a fixed-fee 4-phase delivery roadmap.

Search Queries & Scope Covered: snowflake to oracle autonomous data warehouse migration | reduce snowflake compute credit sprawl oracle | zero etl oracle 23ai converged database vs databricks | oracle select ai natural language sql

SOURCE LEGACY STACK (STACK X)	MODERNIZED TARGET ARCHITECTURE (STACK Y)
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**

ANALYTICS TCO CUT	DATA COPY LATENCY	SECURITY DRIFT	AI TEXT-TO-SQL
45% - 65%	**0 (Zero-ETL)**	**Eliminated (Single VPD)**	**Native Select AI**

1. Problem Space & Architectural Bottlenecks

- Critical business operations remain trapped in High-Cost Snowflake / Databricks Compute Credits, Fragile CDC Copy Pipelines & Egress Fees, creating audit risk, manual reconciliation overhead, and rising support costs.
- Duplicate queries, cloned reports, and undocumented triggers make impact analysis slow and expensive across analytical tables, marts & bi dashboards.
- Lack of automated binary-to-database metadata mapping causes traditional consulting shops to pad estimates with 40% manual discovery buffers.
- Leadership needs a predictable, fixed-fee phased timeline with verifiable test specifications and zero-downtime cutover guarantees.

2. Technical Solution & Reference Architecture

2.1 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.

2.2 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.

2.3 Automated generation of full Business Requirements Documents (BRD), Technical Approach Docs, Testing Specifications, and Verified Test Results.

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

3. Artifact-by-Artifact Conversion & Modernization Matrix

Legacy Artifact / Baseline Workload	Modernized Oracle / Cloud Target	Engineering Method & Outcome
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. Oracle APEX & Ancillary Engineering Toolchain

- Oracle APEX 24.2 & Universal Theme (Light Architecture)
- Oracle Database 23ai / Autonomous Database (ATP / ADW)
- Viscosity Automated Binary (.fmb/.rdf) & AST Metadata Profiler
- Auto-Doc Engine (BRD, Approach Doc, Test Specs & Test Results)
- Oracle REST Data Services (ORDS) & APEX Office Print (AOP)
- SQLcl Liquibase CI/CD, utPLSQL & Playwright Regression Suites

> ***Anonymized Production Proof:*** *Consolidated a fragmented Snowflake and third-party CDC pipeline back into an Oracle Autonomous Data Warehouse (ADW) converged architecture, cutting annual cloud analytics and data-replication spend by 54%.*

5. Principal Architect Q&A

Q: What are the primary levers that control the cost of Snowflake / Databricks to Oracle 23ai ADW?

A: 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.

Q: How does your tool convert binaries (.fmb, .rdf, etc.), map database metadata, and generate documentation?

A: 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.

Q: What does a typical project timeline look like from kickoff to production cutover?

A: 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).