

Transform Netezza workloads to a cloud-native stack or modern data platform – Automatically

Convert legacy data warehouse and ETL workloads 4x faster

Moving legacy Netezza data warehouse and ETL workloads to a cloud-native stack/modern data platform helps enterprises lower licensing costs, release capacity, and enable on-demand scalability.

However, migrating Netezza’s diverse range of artifacts and constructs is a complex task. Netezza scripts contain hundreds of conditional statements (logic), macros, procedural statements, and outputs.

In addition, many enterprises are skeptical of the business and technical risks involved:

- Will there be any business downtime?
- How do I transform years of complex business logic and code?
- Will my workloads be optimized to address the nuances of the new environment?
- How do I prioritize the workloads for transformation?
- How do I seamlessly operationalize Netezza DML scripts, NZPLSQLs, ETL scripts, analytics scripts, and business reports on the target?

The Impetus Workload Transformation Solution addresses all these concerns. With an intelligent grammar engine, it identifies optimization opportunities at schema, code, data, and execution level and automatically converts all types of Netezza workloads, logic, and workflows to target tech stack of your choice.



KEY BENEFITS

1. 4x faster
2. 2x cheaper
3. 4x developer productivity
4. 100% cloud-ready
5. 100% automation across the migration lifecycle



WHAT'S POSSIBLE

- ML-based assessment and recommendation for the target architecture and tech stack
- Business logic conversion
- End-to-end transformation to cloud-native services of your choice
- End-to-end packaging, orchestration, and execution for the target
- Code optimization and query validation to avoid business disruption
- Automated legacy code transformation to a cloud-compatible equivalent with multiple query engine support
- Cost-performance ratio optimization
- Data governance and security compliance

How it works

The Impetus Workload Transformation Solution enables end-to-end automated transformation and operationalization of workloads in four steps.

STEP 1: Assessment and prescription

- Lists entire inventory for diverse workloads
 - Assesses Netezza DML and DDL scripts, data warehouse execution logs, NZPLSQL, shell scripts, scheduler scripts, properties files, and more
- Identifies complex interdependencies to group workloads for offload
 - Provides advanced filters according to workload type and an interactive graphical interface to deep dive into certain flows
- Advanced blueprinting of target architecture
 - Strategizes partitioning, bucketing, clustering, sorting, and distribution keys for improving CPU usage, memory usage, cache hit ratio, and disk I/O
 - Provides comprehensive, configurable recommendations for workload parallelism to ensure optimal performance on target
 - Provides actionable insights and prescriptive recommendations for future-state architecture and tech stack components

STEP 3: Validation

- Pipeline-based automated validation of the transformed code
 - Validates code at the row and cell-level and reports errors with a pluggable tool
- Data-based validation of transformed code
 - Auto-generates sample datasets based on complex query conditions – ideal for unit testing of the transformed queries
 - Feeds the customer-provided dataset for testing on real datasets – suitable for integration testing of the transformed queries

Step 2: Transformation

- Transforms diverse workloads and migrates schema and data to the target store of choice
- End-to-end packaging
 - Transforms core business logic to cloud-native wrappers or orchestrators
 - Converts Netezza SQL grammar, sequences, stored procedures, character sets, and user-defined functions to cloud-native formats like Amazon Redshift, Azure Synapse, or PySpark/PyScala/HiveQL
 - Converts Netezza shell scripts to cloud-compatible shell scripts
 - Re-packages to an open programming language of choice
 - Ensures end-to-end transformation of scheduler/orchestrator scripts to production-ready jobs on target
 - Ensures end-to-end execution on staging and production environment after system integration testing
- Notebook-based inline query editor to optimize or address errors
- Extensible tool and methodology
 - Converts code to a variety of target stores and formats, enabling enterprise-wide use

Step 4: Execution

- Delivers target compatible executable package for cloud-native orchestration
 - Cloud-native orchestration and execution on production
- Ensures optimal performance through parallel execution
 - Provides parallel execution recommendations through exhaustive data-driven assessment
 - Generates required artifacts in the transformation output and executes them in parallel on production
- Supports operationalization
 - Supports end-to-end transitioning into production and operationalization
 - Optimizes target capacity for optimum price-performance ratio
 - Stabilizes environment through a parallel-run period
 - Monitors operations to ensure implicit data governance and compliance on the cloud
 - Ensures continuous integration and delivery
 - Provides runbook documentation, training, and handholding

