IMPETUS | leaplogic™

Automated replatforming of Teradata workloads to Databricks

VIDEO TRANSCRIPT

Are you planning to modernize your data warehouse workloads from Teradata to Databricks, but worried about business disruption?

Enterprises are modernizing their data warehouse workloads to Databricks to leverage cloud scalability, reduce costs, and access advanced analytics.

LeapLogic, Impetus' cloud migration accelerator, can help you achieve this by ensuring an end-to-end automated transformation from Teradata to Databricks.

Here's a demo of how LeapLogic automates and simplifies the transformation of Teradata workloads to Databricks – starting with a comprehensive assessment, followed by automated transformation, validation, and all the way up to operationalization.

LeapLogic first conducts a thorough assessment of your existing Teradata workloads, providing you with valuable data-driven insights. It lists the entire inventory and prioritizes use cases that hold the highest business impact. LeapLogic identifies technical debt within your Teradata environment and provides optimization recommendations for Databricks at schema and orchestration levels.

LeapLogic's comprehensive analysis ensures that your migration journey is meticulously planned and executed. It goes beyond identifying transformation candidates and empowers you with detailed recommendations for each one.

You can go to the Analysis section and view key resource utilization metrics for applications and users to get more details. LeapLogic's schema optimization recommendations are designed to unleash the true potential of your data warehouse. It provides recommendations for Bloom filters, Z-order indexes, partitioning, etc. to improve CPU and memory usage, cache hit ratio, and disk I/O, ensuring your data warehouse operates at peak efficiency. LeapLogic identifies key opportunities for orchestration optimization, streamlining data flow, and eliminating bottlenecks.

Head to the Lineage section to view all interdependencies between various kinds of workloads, along with end-to-end data and process lineage.

In the Recommendations section, you can view suggestions for offloadable workloads and an optimum future-state functional architecture, ensuring cost-efficiency and scalability.

LeapLogic's detailed reports can be downloaded for offline use, ensuring you have access to critical information at your fingertips.

Now, let's explore how LeapLogic automatically converts Teradata workloads to Databricks equivalent.

LeapLogic's Teradata to Databricks pipeline transforms all your data warehouse workloads end-to-end to Databricks. Under the Pipeline section, you can explore various stages of the transformation lifecycle.

The Migration stage converts your Teradata schema to a Databricks-equivalent and seamlessly migrates your data into Databricks tables. The Data Validation stage meticulously validates and certifies the migrated data, ensuring its accuracy and completeness.

The Transformation stage seamlessly transforms your Teradata code and business logic to Databricks-equivalent code. LeapLogic auto-transforms and maps all keywords, functions, and constructs to Databricks-compatible code, ensuring your data warehouse operates efficiently in the cloud.

Presently, Databricks doesn't support operators like All, Any, or Some. LeapLogic follows a unique solution for dealing with unsupported or non-performant patterns on the target side. Its intelligent pattern-based transformation engine incorporates target-native best practices and converts the code to ensure that it is optimized, compliant, and easy to maintain.

LeapLogic also provides a Notebook-based inline query editor for further optimization.

Here, you can see that all transformed SQL queries and business logic are packaged as production-ready jobs along with the orchestration and execution logic.

That's it! You can now download the converted package from LeapLogic containing the target-equivalent code.

The Query Validation stage in the Output section validates the transformed artifacts. As you can see, these transformed queries have been successfully unit-tested on Databricks. The non-unit tested queries require manual intervention to execute successfully on the Databricks platform. LeapLogic's unit and integration testing ensure that your transformed code is rigorously validated before productionization.

Next, you can simply log-in to your Databricks portal and import the transformed code. Here, you can see that the legacy constructs, functions, and keywords have been transformed into Databricks-native equivalent, which can optimally run on the Databricks platform. You can execute the transformed code and notice that multiple Spark jobs have been successfully executed on Databricks.

LeapLogic streamlines your migration journey, ensuring a successful transformation of legacy Teradata workloads to Databricks.

Explore LeapLogic's automation capabilities for the end-to-end transformation of data warehouse, ETL, Hadoop, analytics, and reporting systems to Databricks – faster, at a lower cost, and with minimal risk.

It's more than the next step. It's a leap into the future of your business.