

Migration Assessment and Strategy Services for RDBMS to Apache Cassandra

Assessing, Strategizing, and Accelerating your RDBMS migration to an Apache Cassandra® or DataStax Database

Overview

DataStax provides you a hybrid, multi-cloud, multi-region database with the freedom to run any workload on any device and in any cloud at global scale.

This migration package forms the start of your transformational journey enabling you to ultimately;

- Offload Monolith RDBMS Workload (Oracle and DB2 are supported today)
- Modernize to Microservices Architecture
- Potentially Eliminate Legacy RDBMS
- Improve Self-Managed Operating Economics
- Reduce Friction for Developers and Operators
- Support Advanced Workloads

The constraint that holds back RDBMS-to-cloud transformation projects in enterprises today is the existing application dependency on the data layer for data consistency. This is most easily seen in an application's stored procedures, triggers, and the transactional SQL that the application uses.

Approach

There are several phases to building a proper migration plan to transform an application, with each producing an outcome that feeds into subsequent phases. Workshops and Interviews guide Enterprises through the following stages:

Investigation

Determine the database topology, size, design, schema, utilization patterns, query patterns, and dependencies of the application to be transformed.

Identification

Identify the function, purpose, and use of each data structure. Group data structures and elements that work and function together and help answer some of the following questions: What tables are often selected together with a join? What tables are often used together in a stored procedure? What enforced constraints exist between database structural elements?

Objectification

Determine a persistence and object model for each functional group identified through the prior processes, yielding a new distributed schema design for data persistence.

Functional Isolation

Identify the interface between the application and the data layer. Every bit of logic and code that is below that interface data layer must be incorporated into its appropriate microservice, or the new transformed application API layer. This code can be found in queries, the application itself, or stored procedures.

Occasionally data and functionality will need to be duplicated between application services. Documenting and understanding duplication patterns allows distributed applications to be architected more efficiently.

Transactional Process Identification and Documentation

The transactional elements of the application that have been managed by the monolithic data layer must be identified and documented. The core technical essence of digital transformation is taking transactions that have been managed at the database level and moving them to the application service layer.

Operational Migration Process Definition

With legacy data structures understood and transformed and the new data elements designed and documented, it is possible to plan the operational migration process.

This involves planning to physically remove responsibility from the legacy systems and transferring application servicing responsibility to the transformed system(s). There are two techniques to be considered:

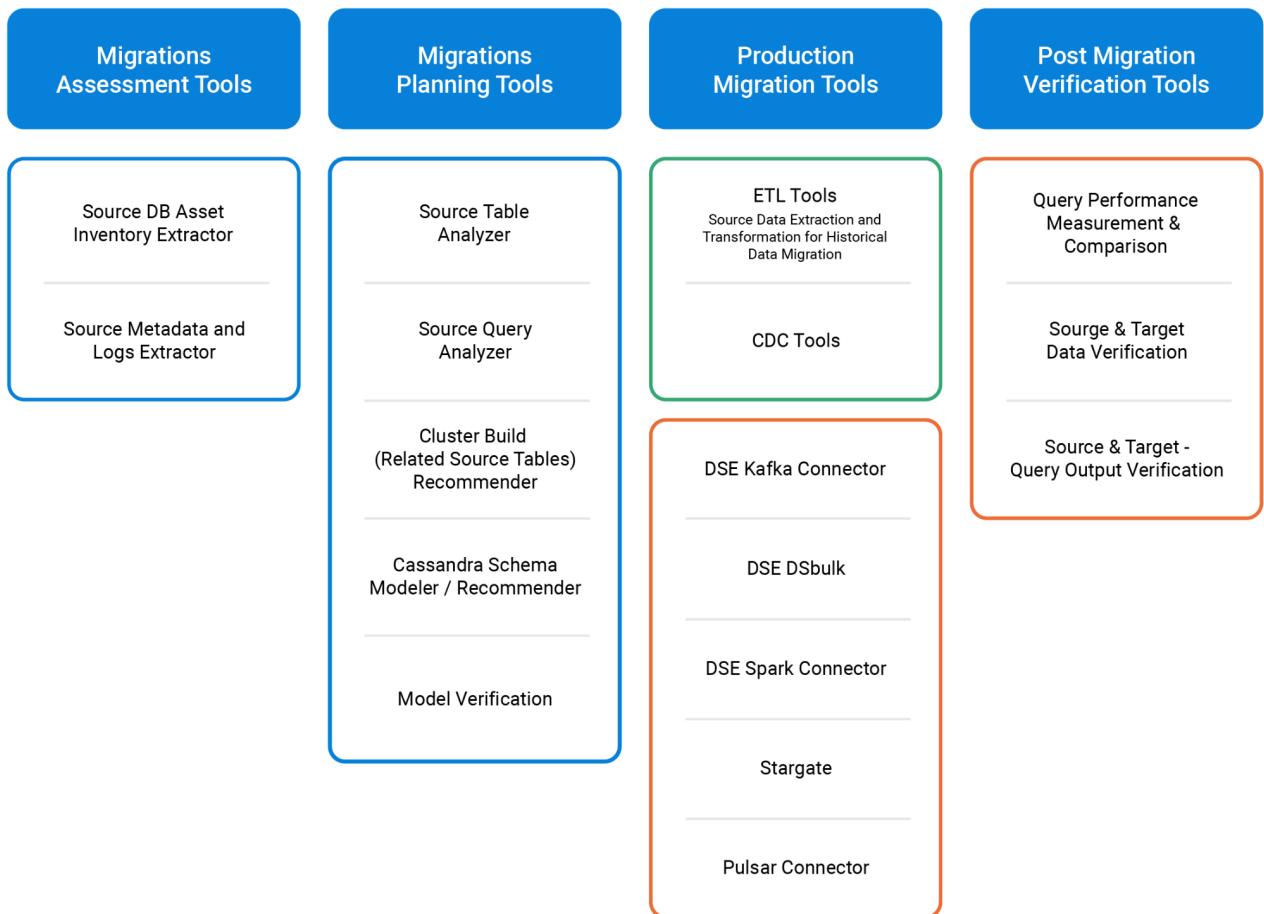
ETL – Extracting the legacy data, transforming it and replicating it as appropriate in a consistent manner to the transformed architecture.

CDC – Change data capture is the process of maintaining consistency between two live systems by capturing the writes to one system, and then transforming and populating the write to another system in near-real-time. This is used to keep legacy and transformed systems consistent throughout live migration processes.

Migration Tooling (with current support for Oracle and DB2)

DataStax has turned the lessons we have learned from experience into automated tooling that makes it easy for developers to solve the hard problems of consistency in the execution of transformation projects.

Datastax offers a robust set of migration accelerators and tools that assist the enterprise with the full lifecycle of the migration effort.



The service package is built around the Migration Assessment and Migration Planning tools of the RDBMS-to-Apache Cassandra Accelerator (blue boxes).

This package focuses on the initial phase of assessment and planning phase of the migration.

1. First, **assessment tools** assist with the challenging task of properly interrogating the original database sources for requirements and query patterns. The discovered patterns are then used to establish collections of tables and data structures that should be migrated together or be grouped into a single service.
2. The **planning tools** then assist the developer with establishing an appropriate object and Cassandra data model for cloud persistence.

RDBMS to Cassandra Migration Planning Scope:

- Kickoff
 - Identify business & IT stakeholders
 - Conduct kickoff with all key stakeholders
 - Plan/organize key discovery workshops to be conducted
 - Define assessment and strategy phase charter, vision and key outcomes
 - Goal: Charter, Workshop schedule and Agenda
- Workshops & Interviews
 - Joint discovery workshops with enterprise development/QA teams and database architects for the purpose of looking into and determining any or all of the following for the purpose of creating the assessment report and proposed migration plans:
 - Investigation
 - Identification using bespoke tooling
 - Objectification
 - Functional Isolation
 - Transactional Process Identification and Documentation
 - Operational Migration Process Definition
 - Migration Methodology Planning
 - *Goal: Discovery & Assessment Reports, Migration Strategy, Recommendations and Proposals, Scope and Plan.*
- Documentation and Training
 - Assessment Reports containing;
 - Options for migration to either DSE, Astra, or Open Source Apache Cassandra
 - Proposal for additional services for PoC, Delivery and TCO for migration option including; migration strategy, migration classification and sizing, scope, schedules, required licenses, tools and services costs

Price

This package is provided on a T&M basis and Customer shall procure 15 Consulting Days. DataStax cannot guarantee that all of the activities/recommendations can be completed during the time allotted, but shall complete as many of the activities/recommendations as time permits.

Terms

- See the [DataStax Services Terms](#) and the [services package description](#) for the details and terms of Services delivery.