

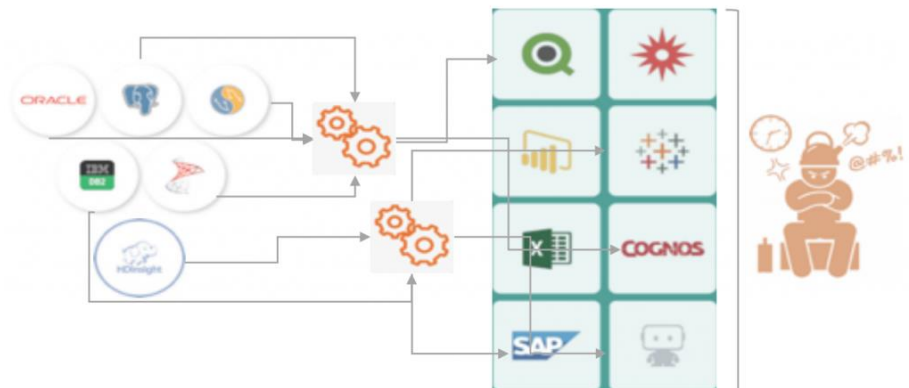


Time & Resource Expensive Queries to Extract Large Data Coming from Multiple Big Data Sources

- 
 - Extracting billions of records
 - Data from multiple data sources
Causing issues with interactive analysis
 - Tedious query performance tuning
 - Redundant Summary Tables
- 
 - Increase query throughput and concurrency
 - Autonomous data engineering approach to manage massive data sets
 - Smart approach to migrate to cloud-based solution
 - Reduce query cost by using aggregations

Challenges

As enterprises turn towards data driven strategies; one of the challenges they face is to have a common business understanding of the data. Different departments in the same enterprises often establish different definitions for the same business processes or data artefacts. For example, what is “net sales”; a department might interpret it as net of invoice line-items and another department might interpret it as net of rebates. Without some level of abstraction or consistency, business is dependent to IT to generate and run reports or risk making big, costly, and worst of all hidden mistakes.



Solution

In order to optimise query performance, enterprises should be looking to eliminate scripting and redundant summary tables. Queries that are constantly scanning raw data should be avoided to increase query throughput and concurrency. Enterprises should invest in define strategies to manage massive data coming from multiple data sources and work with live data for business insights. Having consistent joins between tables and apply Machine Learning (ML) can reduce data query times from 2X to 1000X as well. Enterprises need

a smart approach to embracing the cloud and significantly reduce data warehousing and query costs. Eliminating the need to copy or move data across the infrastructure can also reduce the query cost significantly. With query times shortened from days to hours or minutes, BI teams can deliver time-sensitive analytics to the line of business effectively and efficiently.

As an approach to data management that allows applications to use data without requiring technical details about it, such as how it is structured, where it is physically located, or how the data is accessed.

How AtScale can help

AtScale's autonomous data engineering, as queries are run against datasets in the adaptive analytics fabric (the virtualized data and associated tools to aid analytics speed, accuracy, and ease of use), machine learning is applied to determine what data within the larger set is needed.

AtScale's A3 platform optimizes analytical workloads based on user behaviours, using artificial intelligence. AtScale monitors, analyses, and optimizes data queries based on how data is being used, learning what data is important, and how best to shape that data into aggregates, providing conversational response time.

Because extraneous data is bypassed altogether during the query process, data is delivered between 2x and up to 1000x faster, depending on which data sets are queried. Queries that run against many terabytes of data can yield responses in tenths of a second from an AtScale aggregate.

Collectively, these aggregates are referred to as 'Accelerated Data Structures' because they constantly adapt to changes in user behaviour. With query times shortened from days to hours or minutes, BI teams can deliver time-sensitive analytics to the line of business effectively and efficiently.

