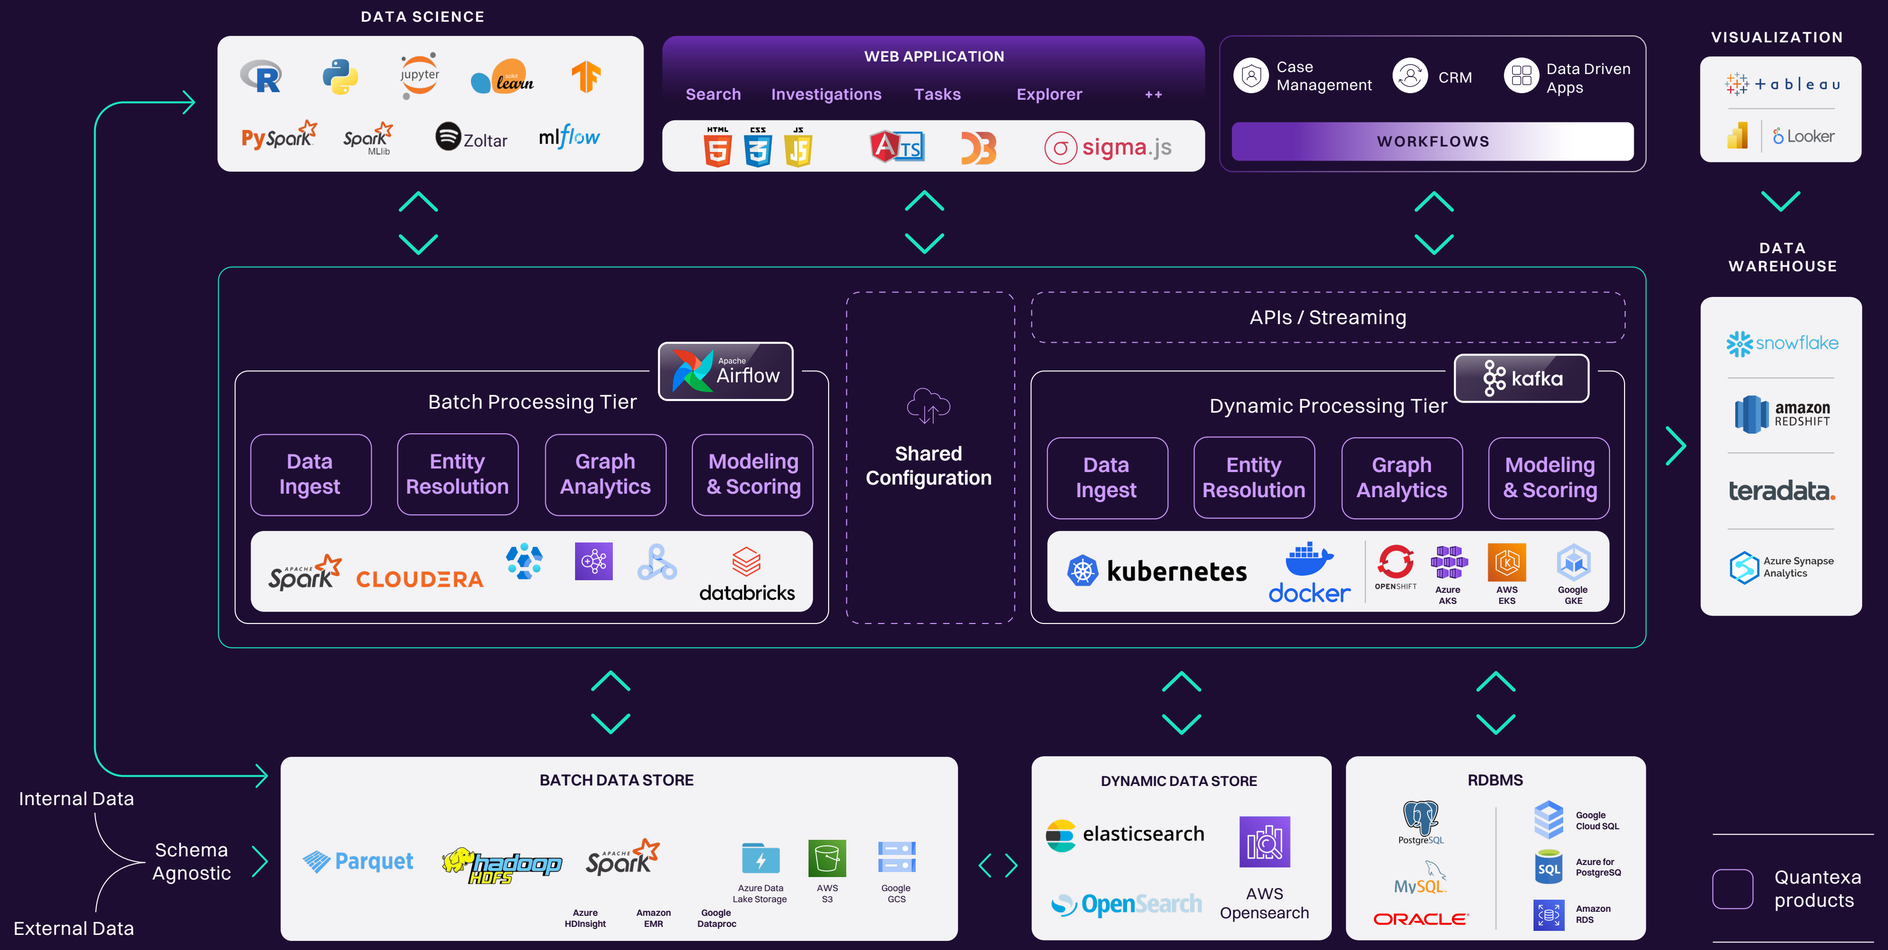
quantexa Decision Intelligence Platform - Reference Architecture

An end-to-end platform that unifies data, uncovers context, and powers human and AI decisions to build a solid data foundation. Our Decision Intelligence Platform can be deployed in on-prem, cloud, or hybrid environments. Our open and extensible architecture makes it easy to get data in and out of the platform with scalable APIs and streamlined integrations with the downstream applications and solutions you use the most.



Data Layer

Batch Data Store: Preferred data lake or distributed file system.

Dynamic Data Store: Indexed data for real-time processing through APIs and streaming.

RDBMS: State for saved investigations, tasks, and audit logs.

Batch Processing Tier
A Spark engine built in native Scala that performs linearly scaling bulk data operations providing easy access for data science, analytics or batch processing for tens of billions of records.

Dynamic Processing Tier
A set of microservices that resolve entities, builds graphs and calculates scores dynamically, providing search, investigation management, and processing streaming updates using Kafka. Builds entities and networks on the fly for multiple use cases without data duplication.

Shared Configuration
Out-of-the-box and configurable Data Ingest, Entity Resolution, Graph Scripting and Scoring for both batch and dynamic processing allowing seamless transition from analytics to real-time operation.

Quantexa.com