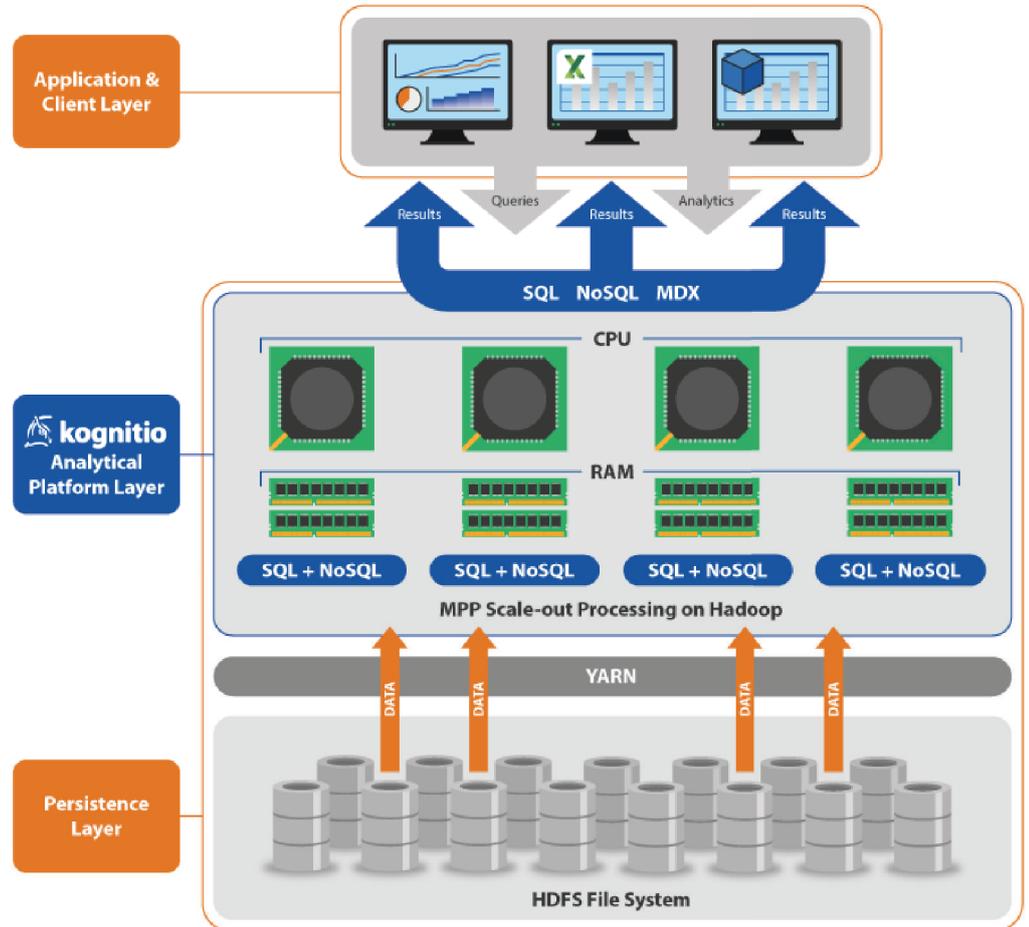


Ultra-fast High-concurrency SQL on Hadoop

Kognitio on Hadoop is the latest version of the Kognitio Analytical Platform, the world's fastest in-memory data analysis engine. Kognitio on Hadoop includes full YARN (Hadoop's preferred resource manager) integration allowing Kognitio to share hardware infrastructure with other Hadoop applications.



Kognitio on Hadoop provides ultra-fast SQL access, at very high concurrency levels, to data stored in Hadoop. The speed and concurrency allows organisation to deploy interactive, self-service analytical tools to thousands of end users even when the data volumes are very large ("Big Data"). Hadoop and its associated software projects have provided organisations with a cost-effective way of storing and processing vast amounts of diverse data and a range of tools that allow that data to be captured analysed and to some degree, consumed. However, mass consumption of Hadoop based datasets by large numbers of users in a business context, is still difficult, as the tools available do not possess the performance levels or enterprise capabilities needed to reliably support hundreds of concurrent users. Kognitio on Hadoop can support thousands of analytical queries per second from hundreds of concurrent sessions and can be scaled out to accommodate both larger data volumes and more query throughput. In addition to the ultra-fast SQL Kognitio also has sophisticated NoSQL support that allows advanced analytical algorithms to be executed at scale and in almost any language e.g R, Python. In-line with the rest of the Hadoop market, Kognitio on Hadoop is available free of charge with no restrictions on size or functionality.

How does it do it?

Data is held in very fast high-speed computer memory (RAM)

The architecture is shared nothing MPP; each CPU core operates on its own individual chunk of memory; this "shared nothing" approach has been a Kognitio hallmark from the software's earliest days

Data is held in structures that are optimized for in-memory analysis; it is not a transient copy of disk-based data i.e. like a traditional cache

Massively Parallel Processing (MPP) allows platforms to be scaled-out across the largest of Hadoop clusters

True query parallelization allows queries on very large data-sets to equally use every processor core, on every processor (CPU), on every server

Granular queries that access smaller sub-sets of data can be automatically isolated to a sub-set of CPU cores thus allowing hundreds of these queries to be satisfied simultaneously with zero computational contention

Processor efficiency is very high. Kognitio uses development languages and sophisticated techniques to ensure every CPU cycle is effectively used

Machine code generation and advanced query plan optimization techniques further ensure every processor cycle is effectively used to its maximum capacity

About Kognitio

For more than a generation, Kognitio has been a pioneer in software for advanced analytics, helping companies gain greater insight from large and complex volumes of data with low latency and limitless scalability for competitive business advantage. Sitting at the nexus of Big Data, in-memory analytics and cloud computing, Kognitio extends existing data, BI and visualization investments as an analytical accelerator, providing a foundation for the growing community of business analysts, self-serve users and data scientists. The Kognitio Analytical Platform can be used as a powerful tool in a data science lab or to satisfy mass scale queries from hundreds of demanding users; it runs on industry-standard servers, as an appliance, on Hadoop or in the cloud managed by the Kognitio Analytical Services team. To learn more, visit kognitio.com and follow us on LinkedIn, Twitter and Facebook



facebook.com/kognitio



linkedin.com/company/kognitio



twitter.com/kognitio



youtube.com/kognitio



info@kognitio.com