Enterprise-grade Big Data Analytics with SAP HANA Vora™ and MapR

SAP HANA Vora at a Glance

SAP HANA Vora is a distributed computing solution for business that leverages and extends the Apache Spark execution framework to provide enriched interactive analytics on Hadoop.

Product Snapshot

In-memory query engine running on Apache Spark execution framework

Compiled queries for accelerated processing across nodes

Enhanced Spark SQL semantics with hierarchies for analytical processing

Enhanced mashup application programming interface (API) for easier access to enterprise application data for machine learning workloads

Apache Spark-based OLAP on Hadoop with SAP HANA Vora and MapR
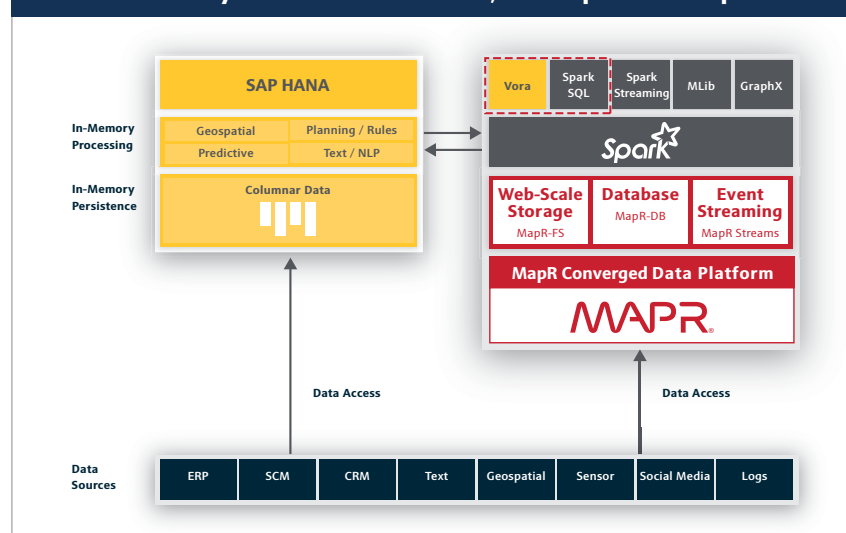
OLAP analysis techniques have existed for many years and enabled analyses such as customer profitability, product profitability, and other types of analysis. The explosion in the volume and variety of data has led to the fast adoption of Hadoop for storage, archival, and processing of large amounts of data. There is an increasing need to deliver analysis on this new data that's originating from sources such as mobile, social, and IoT, but also tie it back to structured data sources such as ERP systems. This will enable enterprises to get contextual insights and enable precision decision-making. This is where OLAP analysis on big data comes into the equation.

SAP and MapR have partnered to ensure customers can perform interactive and drill-down analysis across corporate and Hadoop data. The combination of MapR and SAP helps bridge the divide between enterprise data and big data, with simplified ownership of big data across IT. The MapR Platform integrates natively with SAP HANA Vora with Spark as the execution engine that enables access to data stored in the MapR Platform.

Realize greater than 5X ROI with the MapR Converged Data Platform

MapR provides the industry's only Converged Data Platform that integrates the power of Hadoop and Spark with global event streaming, real-time database capabilities, and enterprise storage, enabling customers to harness the enormous power of their data. MapR enterprise-grade features include NFS for real-time data ingest, snapshots and mirroring capabilities for true high availability and disaster recovery, and multi-tenancy support ensuring secure access to data. The MapR File System is a fully read-write POSIX compliant file system with an NFS and HDFS API.

Contextual Analytics across SAP HANA, Hadoop and Enterprise Data

Benefits of SAP HANA Vora™ and MapR

Get Started with MapR and SAP Today!

Get the MapR Sandbox for Hadoop, a fully functional Hadoop cluster running on a virtual machine. Visit mapr.com/sandbox

Get the SAP HANA Vora developer edition on AWS here go.sap.com/cmp/syb/crm-xu15-int-voratrdm/index.html

Bridge the divide between enterprise data and big data

SAP HANA Vora solves key big data challenges by providing:

Data correlation for making precise contextual decisions Enables mashup of operational business data with external unstructured data sources for more powerful analytics.

Simplified management of big data Allows data to be processed locally on a Hadoop cluster, removing any data ownership and integration challenges.

OLAP modeling capabilities on Hadoop data. Real-time drill-down analysis is possible on large volumes of Hadoop data distributed across thousands of nodes.

SAP HANA Vora enables OLAP analysis of Hadoop data through data hierarchy enhancements in Spark SQL and compiled queries for accelerated processing. SAP HANA Vora helps make decisions more precise through interactive and drill-down analysis across corporate and Hadoop data. Data scientists and developers can now readily enrich their dataset in Hadoop with corporate data in the SAP HANA® database. It also reduces complexity and cost by simplifying data access and processing across SAP HANA and Hadoop.

Why MapR and SAP HANA Vora

MapR and SAP HANA Vora have partnered to deliver OLAP based Hadoop solutions with enterprise-grade capabilities.

• Easily scalable with the ability to go to thousands of nodes
• Built-in High Availability and Disaster Recovery via the MapR Platform
• Ability to run enhanced Spark SQL semantics seamlessly on the MapR Platform to include hierarchies that enable OLAP and drill-down analysis
• Lower the cost of keeping archived data highly accessible

About SAP

As market leader in enterprise application software, SAP (NYSE: SAP) helps companies of all sizes and industries run better. From back office to boardroom, warehouse to storefront, desktop to mobile device—SAP empowers people and organizations to work together more efficiently and use business insight more effectively to stay ahead of the competition. SAP applications and services enable more than 291,000 customers to operate profitably, adapt continuously, and grow sustainably. For more information, visit sap.com.

About MapR

MapR provides the industry's only Converged Data Platform that integrates the power of Hadoop and Spark with global event streaming, real-time database capabilities, and enterprise storage, enabling customers to harness the enormous power of their data. A majority of customers achieves payback in fewer than 12 months and realizes greater than 5X ROI. MapR ensures customer success through world-class professional services and with free on-demand training. Amazon, Cisco, Google, HP, Microsoft, SAP, and Teradata are part of the worldwide MapR partner ecosystem. Investors include Google Capital, Lightspeed Venture Partners, Mayfield Fund, NEA, Qualcomm Ventures and Redpoint Ventures.