

MapR-DB is a fast, scalable, and enterprise-ready in-Hadoop database architected to manage big data. It uses a wide-column (or “table-style”) NoSQL data model to support millions of columns across trillions of rows—up to one trillion tables. It leverages the Apache® HBase™ API so existing HBase applications can be migrated easily to MapR-DB. Designed to handle business-critical operations on unstructured and non-relational data, it is proven in production at Global 2000 companies across many industries, as well as at Web 2.0 companies.

MapR-DB is well suited for deployments requiring operational analytics, in which analysis is performed on live operational data. Example operational analytics use cases include real-time fraud detection, security information and event management (SIEM), ad targeting, and recommendation engines. MapR-DB is integrated with the MapR Distribution including Apache Hadoop, which means customers can run NoSQL operations directly on the top-ranked Hadoop distribution, in a single consolidated cluster. This simplified architecture eliminates the slow and bandwidth-intensive data copying required in multi-cluster environments. Another advantage is low-latency analytics, including real-time processing of incoming data feeds. The architecture also eliminates duplication of data and resources, and reduces the overall data center footprint for lower total cost of ownership.