

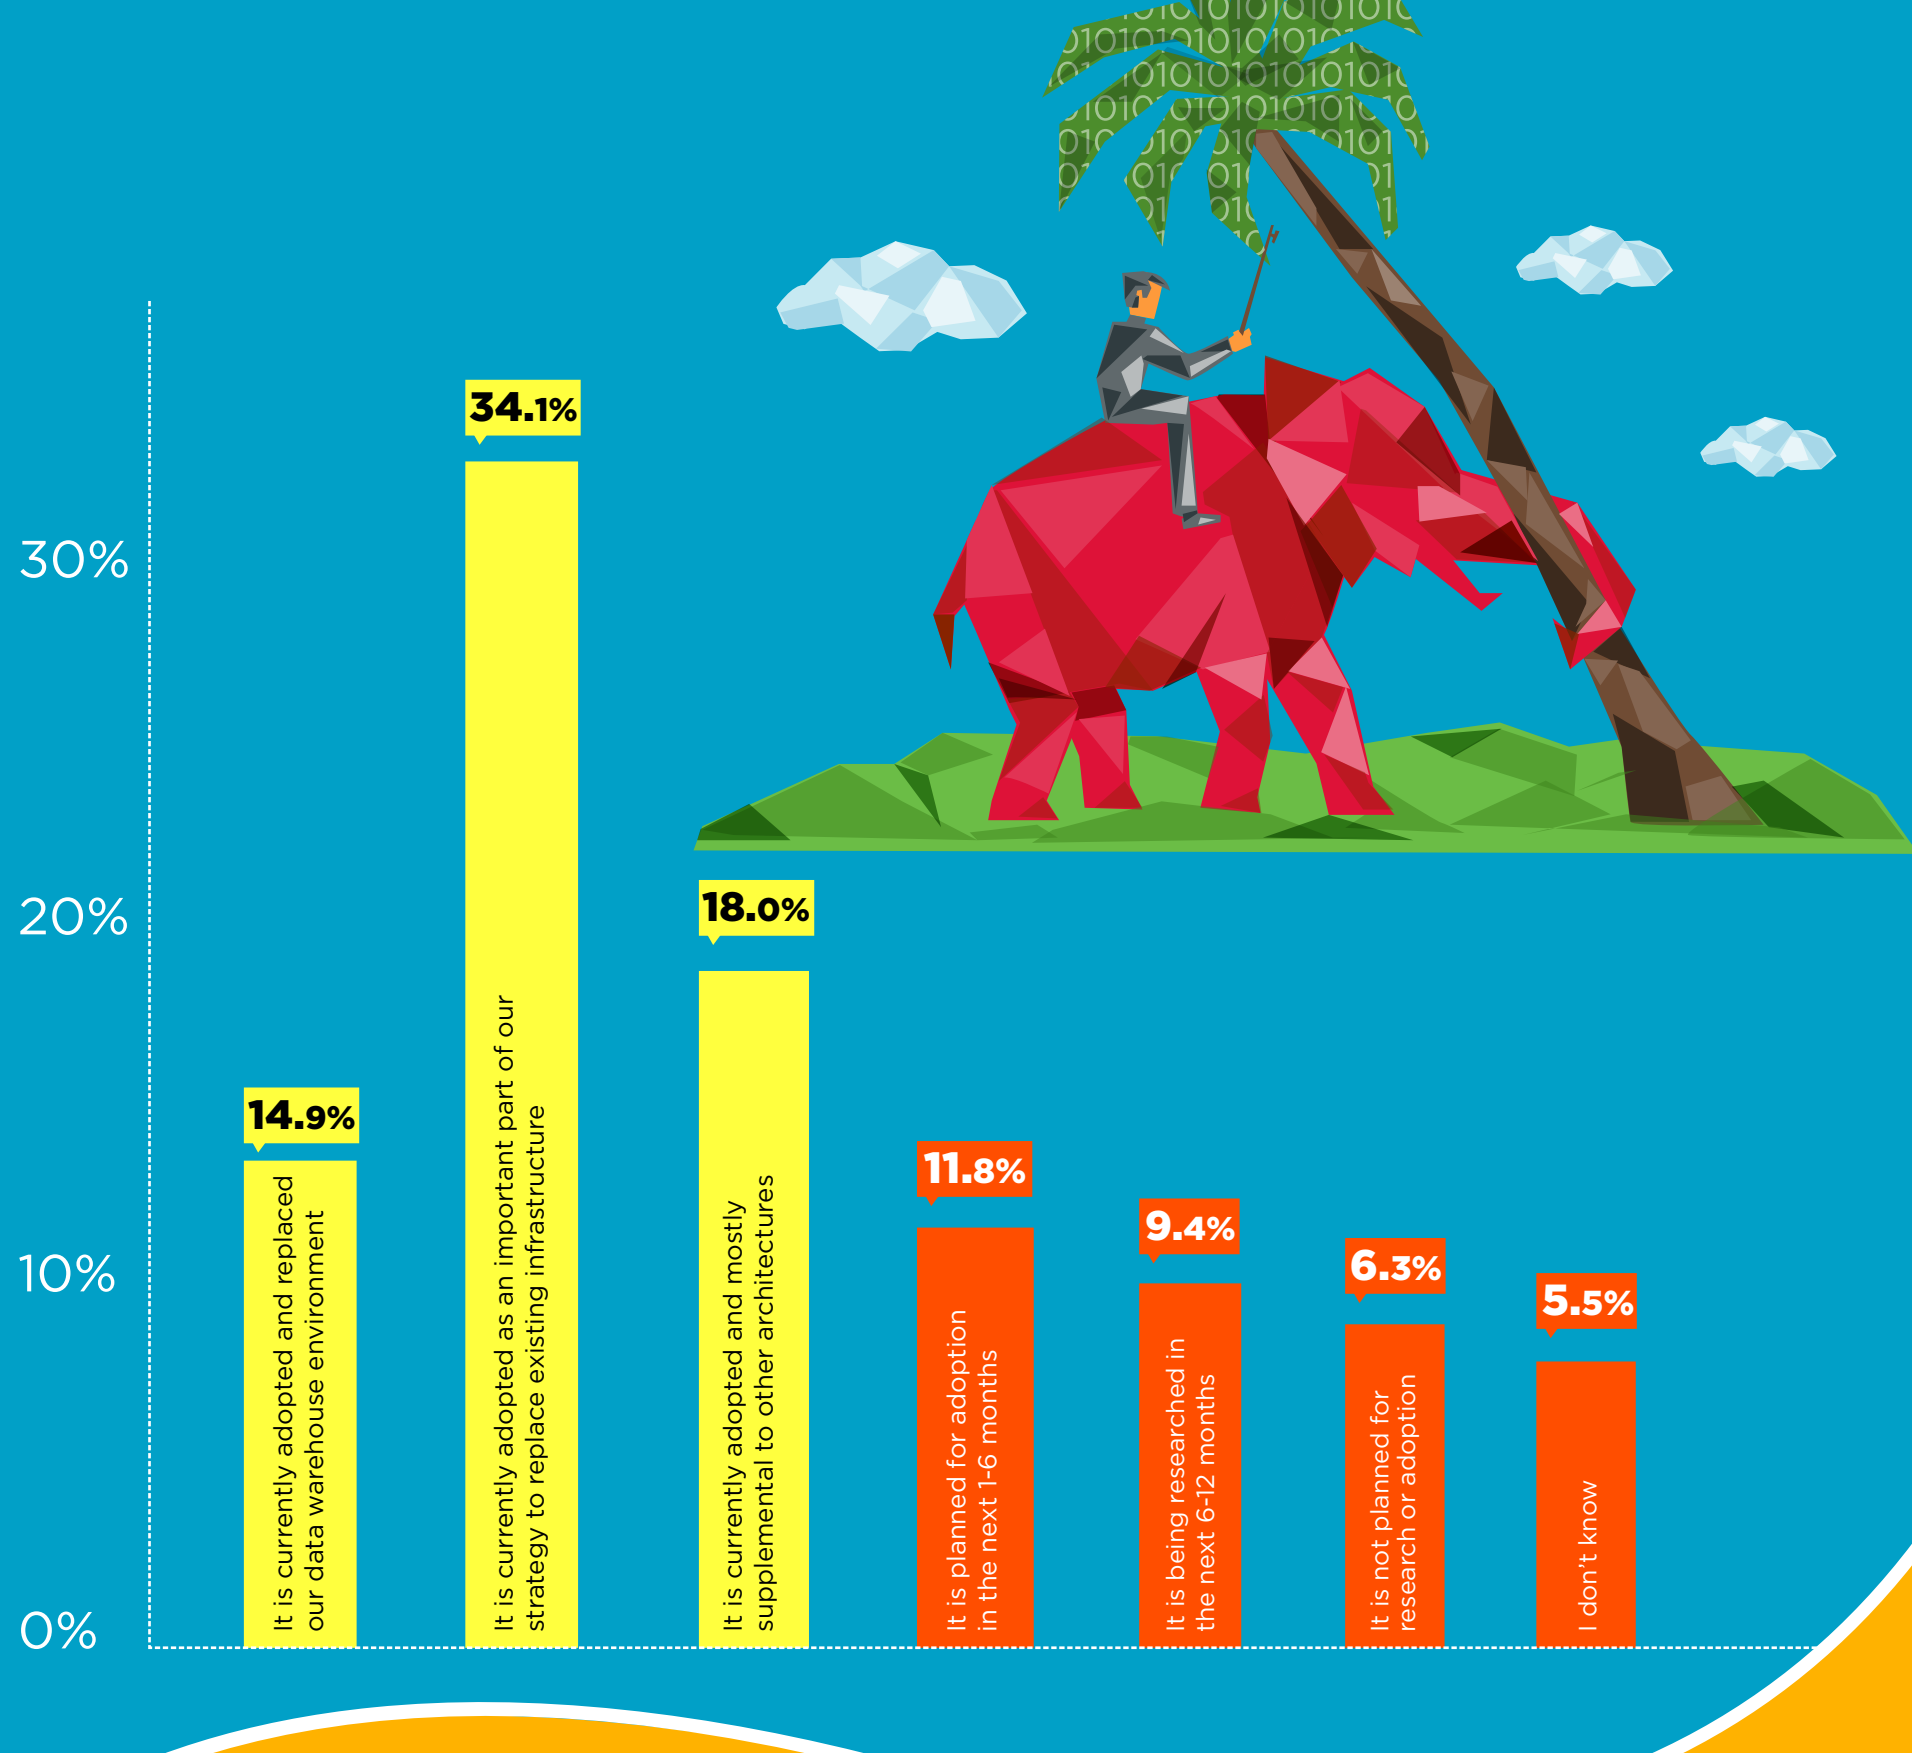
HOW DOES Hadoop Help Elevate Your Data Warehouse?

Enterprise data warehouses (EDWs) are growing in size and complexity due to rapidly changing data requirements. The increase in data complexity and heightened demand for access from business stakeholders places a burden on EDWs to support increasing numbers of data consumers.

New data sources from IoT sensors, mobile and online apps, and machine-generated sources such as web and application logs create these new data requirements. These requirements create a burden that EDWs were simply not designed to handle.

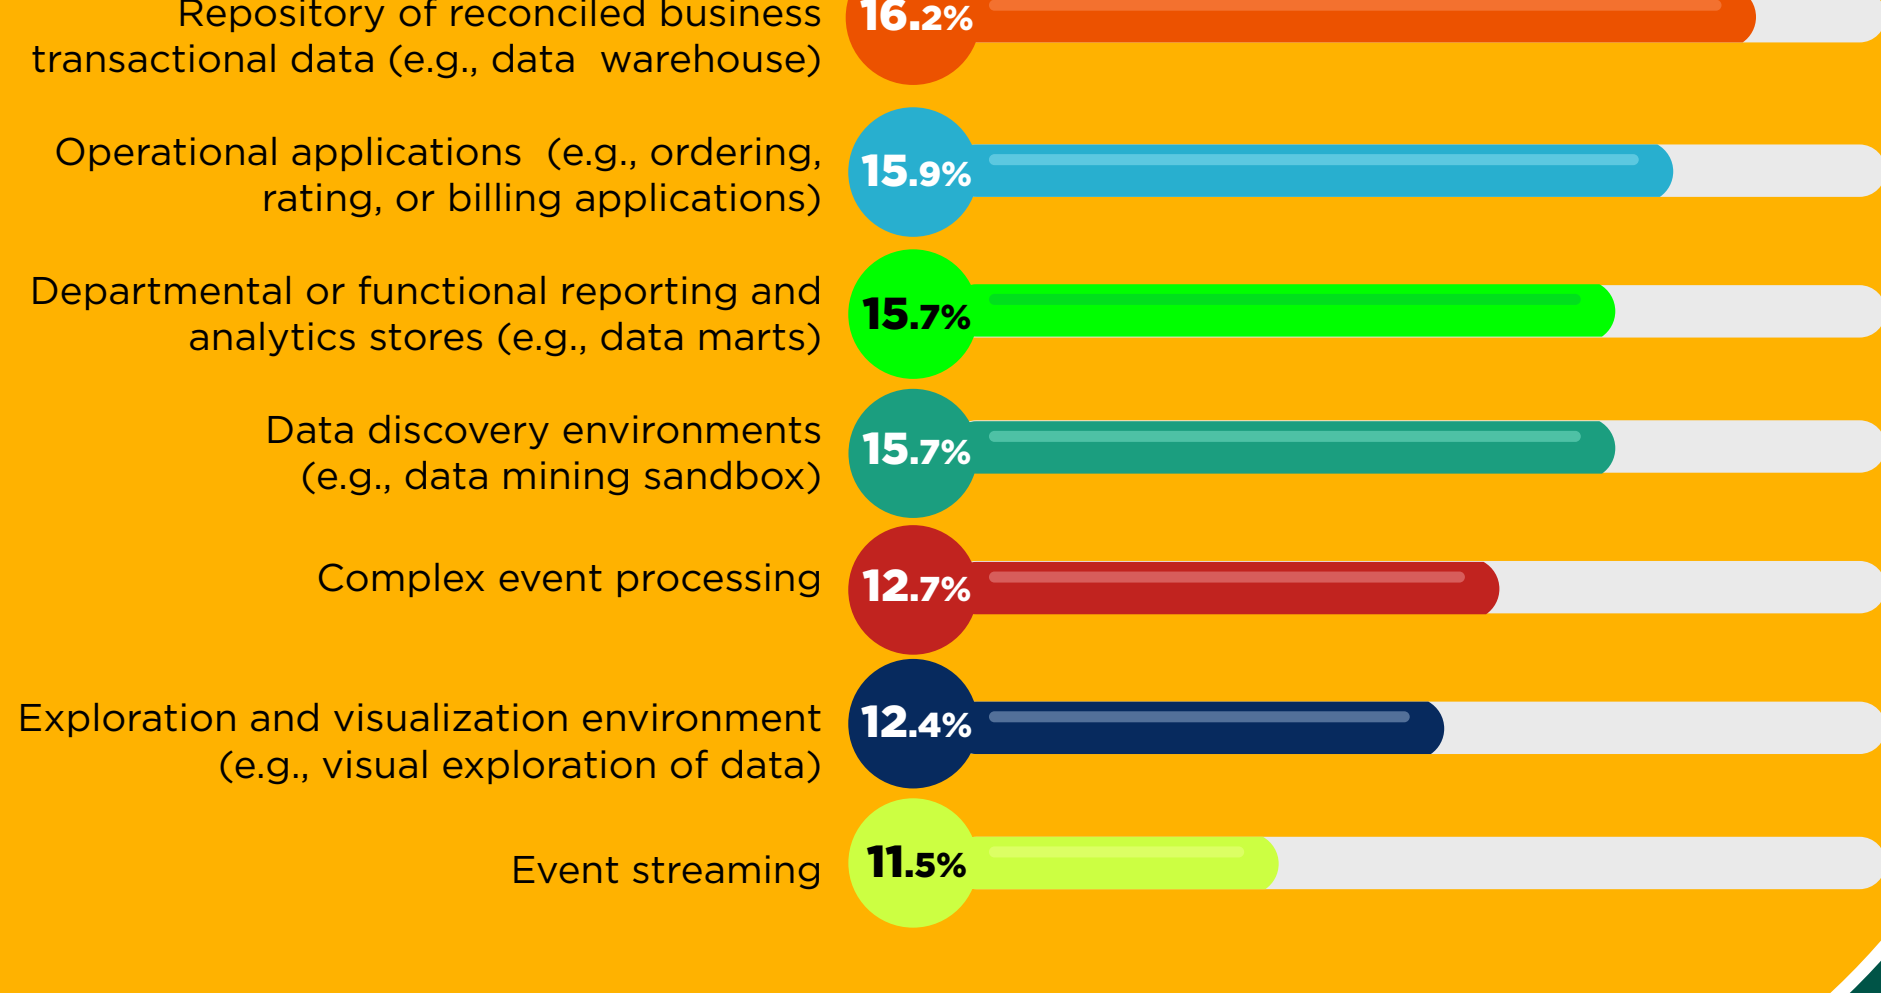
ENTER HADOOP AND THE DATA LAKE

Data lakes based on Hadoop offer economical alternatives for data ill-suited for the EDW. Recent EMA research found over **2/3** of companies have a data lake strategy.



EVOLUTION OF THE DATA LAKE

The data lake is now more sophisticated. It can support multiple analytical use cases and workloads, supporting many of the functions of the EDW.



NEED FOR SQL ON HADOOP FOR BUSINESS PROJECTS

To increase adoption and distribution of these projects, data lakes need to support SQL-based tools to provide business stakeholders with direct access.



EASE THE LOAD OF THE EDW

Reduce the burden on your EDW by using a data lake strategy. This strategy includes your EDW **AND** Hadoop to support new data complexities and provides:



An economical platform for storage



Support for the processing of multiple workloads to meet the evolving requirements of the data lake



The ability to meet the demand of a robust SQL interface to provide business stakeholders with direct access to the data to make decisions about customers, revenues, and daily operations

VERTICA ON HADOOP

MEETS THE CHALLENGE TO ELEVATE YOUR DATA WAREHOUSE

Vertica SQL on Apache Hadoop® offers an enterprise-ready method to perform SQL queries on your Hadoop data. Leveraging years of experience in the big data analytics marketplace and opening Vertica to use the full power of Hadoop, organizations can perform analytics regardless of the data format. Vertica SQL on Apache Hadoop handles your mission-critical analytics projects by merging the best of their analytics platform with the best that Hadoop data analytics can offer. The principles below help to deliver on the promise of:




A unified analytics platform merging the economics of Hadoop with the power of Vertica



Complete SQL support to empower business stakeholders to access and use data directly



Using Orc and Parquet files on HDFS for Information Lifecycle Management



Workload management to support scheduled and ad-hoc analytics workloads from across the organization

Advanced Analytics

Gain full-functionality ANSI SQL capability, not a subset of commands. Run 100% of TPC-DC benchmark queries with no modification.

Enterprise Ready

Install Vertica directly on your Apache Hadoop infrastructure with no helper nodes or single point of failure.

Platform Agnostic

Integrated with any distribution of Hadoop, including Hortonworks, Cloudera, and MapR.

Data Access

Data query across Parquet, ORC, JSON, and many other formats.

For more information, visit:

<http://www.vertica.com/SQLonHadoop>