



# Streaming Fast Data Applications into Hadoop with VoltDB Export



Ideally suited as an operational database to rapidly ingest and process streams of fast data, VoltDB has the ability to export data at high speed to long-term Hadoop-based analytics stores via VoltDB Export.

VoltDB Export transactionally pushes data from VoltDB into Hadoop using the Web HDFS HTTP connector. Unlike an Extract-Transform-Load (ETL) process, VoltDB Export uses a push model that exports data at the same rate at which it is ingested, avoiding slowdowns and eliminating overflows to disk under high throughput.

VoltDB allows application developers to automate the export process by specifying certain tables in the schema as sources for export. At runtime, any data written to the specified tables is sent to the export connector. The Hadoop HTTP export connector receives serialized data from the export tables and writes it out to Hadoop via HTTP requests to WebHDFS.

The VoltDB Export system is a loose coupling managed from within the VoltDB application. The application has complete control over when and what data moves to the external system. The Hadoop export connector is managed by the database servers, which distribute the work to ensure maximum throughput.

For application developers and enterprises looking for an end-to-end solution that combines the long-term, deep analytics capability of Hadoop with the speed and operational capability an in-memory analytics solution, the Export delivers the operational capabilities of VoltDB (ingest, enrich, analyze, decide) while enabling data to be moved out of VoltDB once the data has been processed and is ready for longer-term analysis and storage.

## VoltDB Export and Hadoop

### Fast Ingest

VoltDB's rapid ingest capability enables hundreds of thousands to millions of data events to stream in to the system at wire speed for enrichment and real-time analysis.

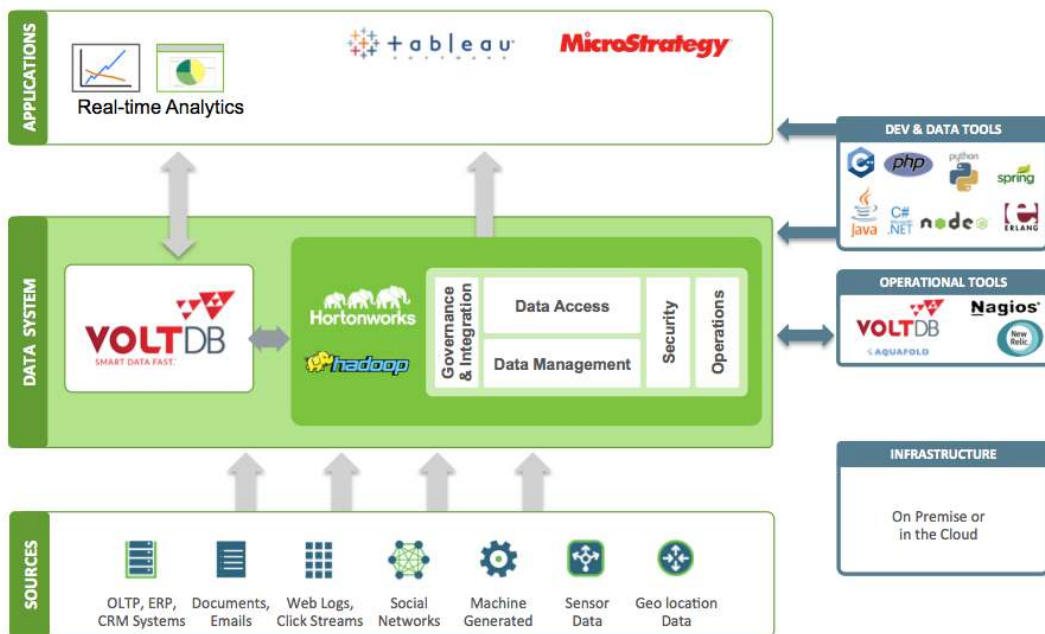
### Real-time Analytics inform real-time decisions

VoltDB's real-time analytics processes millions of in-coming data events, providing insight to power real-time decision-making.

### Reporting and long-term analytics

Processed data from VoltDB is exported at the speed at which it was ingested to Hadoop for long-term analytics, reporting, data science and storage.

# VoltDB in the Modern Data Architecture



## VoltDB is:

An in-memory, scale-out, SQL operational database for high-performance, fast data applications

- Write and read millions of events per second
- Export processed data to Hadoop as fast as it comes in
- Make automated, per-event decisions based on current and historical data

## VoltDB and Hadoop

VoltDB enables developers to take advantage of the cyclical nature of the import-export data cycle. For example, inflows of data can be filtered and acted upon based on rules loaded into VoltDB. Based on this filtered and processed export stream, updated rules are generated in Hadoop and frequently reloaded into VoltDB. VoltDB Export enables data to arrive in your analytic store sooner, and allows deep analytics to be leveraged with radically lower latency.

## VoltDB is a Certified Technology Partner

The Hortonworks Certified Technology Program reviews and certifies technologies for architectural best practices, validated against a comprehensive suite of integration test cases, benchmarked for scale under varied workloads and comprehensively documented.



## Hortonworks. We do Hadoop.

Hortonworks is a leading commercial vendor of Apache Hadoop, the open source platform for storing, managing and analyzing Big Data. Hortonworks Data Platform, our distribution of Apache Hadoop, provides an open and stable foundation for enterprises and a growing ecosystem to build and deploy Big Data solutions.

Hortonworks is the trusted source for information on Hadoop, and together with the Apache community, Hortonworks is making Hadoop an enterprise data platform. Hortonworks provides unmatched technical support, training and certification programs for enterprises, systems integrators and technology vendors.

For additional questions, contact:

- VoltDB
- <http://voltdb.com>  
(978) 528-4660
- Hortonworks  
[www.hortonworks.com](http://www.hortonworks.com)  
(855) 8-HORTON