

eBOOK

HOW VOLTDB ENABLES EVENT-DRIVEN DECISIONS FROM KAFKA



VOLTDB

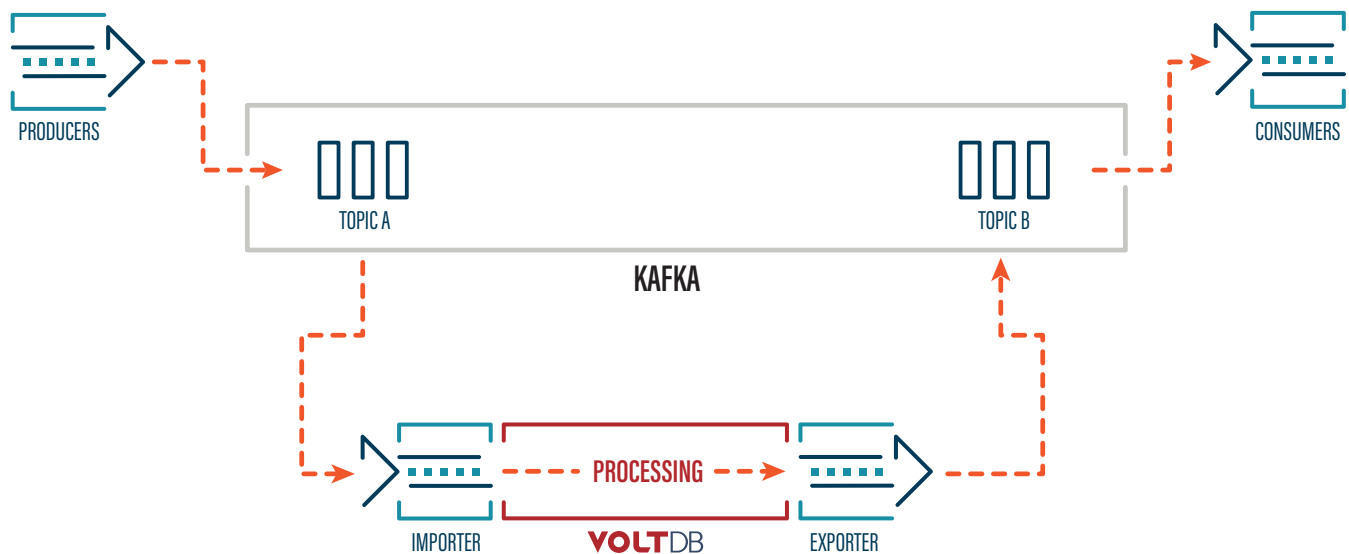
INTRODUCTION

VoltDB is a high-velocity decisioning engine, powering applications that must react in milliseconds— rather than seconds or minutes—in order to drive revenue or prevent revenue loss. Examples include BSS (policy and charging), fraud prevention, customer value management (i.e. hyper-personalized offers), and real-time industrial automation.

The nature of these applications often requires VoltDB to operate within a diverse and heterogeneous computing ecosystem where it needs to “play well” with other technologies, including Apache Kafka. Kafka has established itself as the de-facto enterprise message queue choice. Kafka’s rich ecosystem of Kafka Connectors further eases the integration with the other technologies that are part of the architecture already.

While there are several choices to siphon data from Kafka for further processing downstream, when it comes to being able to process data to drive event-time in-stream decisions, choices become sparse.

VoltDB ingests data directly from Kafka by subscribing to Kafka topics, manages ingest-through-decision in under 10 milliseconds, then sends the decision back to Kafka so an action can be taken immediately, before fraudsters have a chance to do damage or a monetization opportunity is lost.



VoltDB adds value to your Kafka investment by complementing Kafka technology in precisely the way you need it to: by enabling real-time decisions on complex streaming data and allowing your applications to focus on the important work of analyzing, processing, and modifying data in flight through secure, reliable transactions.

Our customers combine VoltDB with Kafka to power diverse use-cases, ranging from a mobile telephone operator implementing real-time personalization for customers in over 90 countries to a large sports fantasy platform needing highly concurrent performance at scale to deliver an exceptional user experience.

You've already made the investments you need to make to create a mature, event-driven architecture. Why short yourself by cutting corners and not getting the most that you can out of those investments by not combining them with technology that enables them for the current age of 5G, IoT, and machine learning?

Given the rise of 5G, machine learning, and IoT, VoltDB has become a natural technology complement to Kafka, enabling it to quickly and seamlessly query complex streaming data. Kafka was built (in 2011) to process streaming data but it was built long before big data truly became fast data via 5G, IoT, and machine learning. Hence, while Kafka works okay for some use cases, it won't work well for use cases that require rapid querying of complex streaming data.

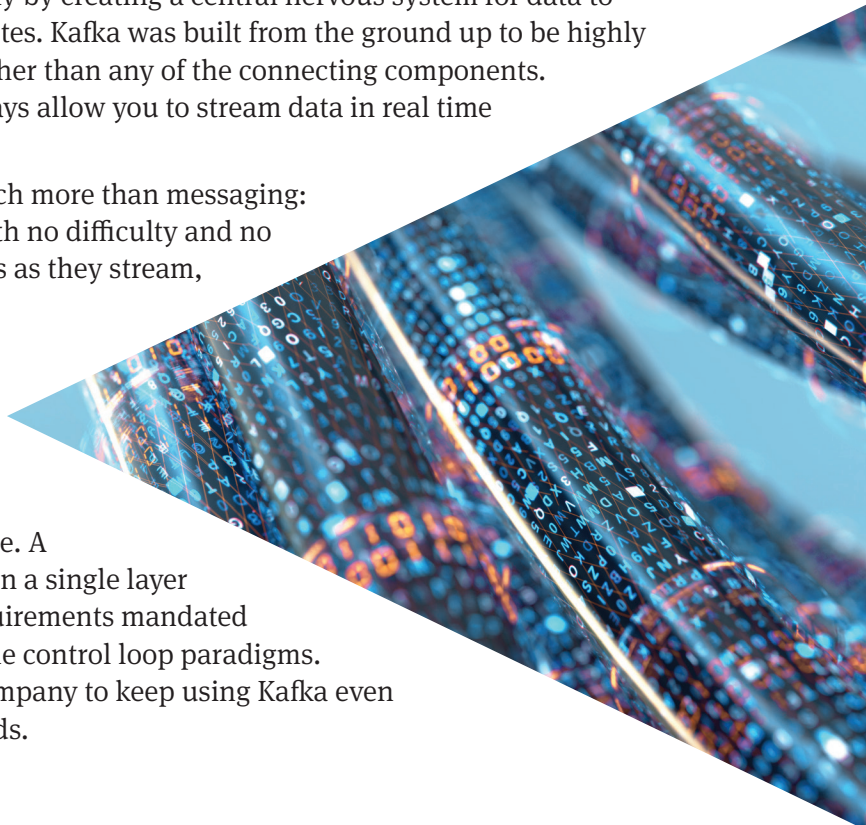
THE POWER OF KAFKA

Today's economy is based on knowledge and data. As in the oil economy of yesterday, companies need to build an infrastructure to process this knowledge and data efficiently and effectively.

Apache Kafka reduces this complexity dramatically by creating a central nervous system for data to flow through, creating nothing but linear data routes. Kafka was built from the ground up to be highly available with throughput that is significantly higher than any of the connecting components. Lastly, it was built as a streaming platform to always allow you to stream data in real time with no issue.

Kafka isn't just messaging done right, it's also much more than messaging: it's unlimited scalability and easy data storage with no difficulty and no performance penalties from storing your messages as they stream, and it's all done in real time.

However, while Kafka and KsqlDB together can make these streams readable using SQL-like language (KQL), to store and process the data meaningfully one would need multiple auxiliary technologies like a NoSQL database, a stream processing technology, and a rules engine. A combination of these capabilities need to coexist in a single layer to meet the stringent latency and intelligence requirements mandated by 5G, Industrial IoT, and their associated real-time control loop paradigms. That's where VoltDB comes in. We enable your company to keep using Kafka even with these incredible new data volumes and speeds.



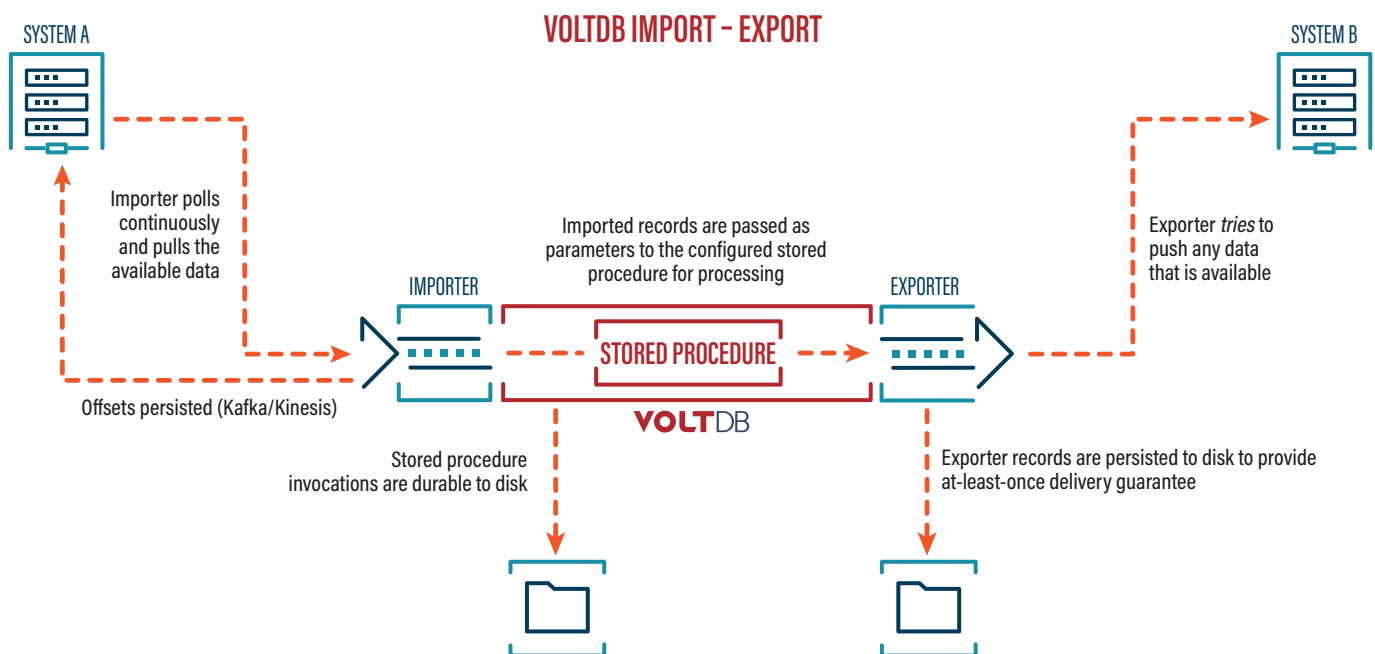
YOUR CHOICE: INTEGRATE OR REPLACE

Depending on the architecture and the current point in the journey towards an event-driven architecture, applications can integrate VoltDB with a Kafka message queue by using the importer-exporter framework or can use VoltDB as a Kafka message queue by itself using the new Topics feature.

Here's a brief overview of both options.

Integration using Importer-Exporter framework

VoltDB's importer-exporter framework provides a versatile no-code integration with other technologies like Kafka, Kinesis, JDBC, Elastic, Hadoop, etc. Importers stream data into VoltDB for per-event ingestion while ensuring the durability of the messages/records after they are pulled into VoltDB. Exporters can then push the post-processing messages/records to a downstream system with an at-least-once guarantee.



Import

The Importer framework enables applications to integrate VoltDB into an event-driven architecture. A VoltDB Importer manages the process of:

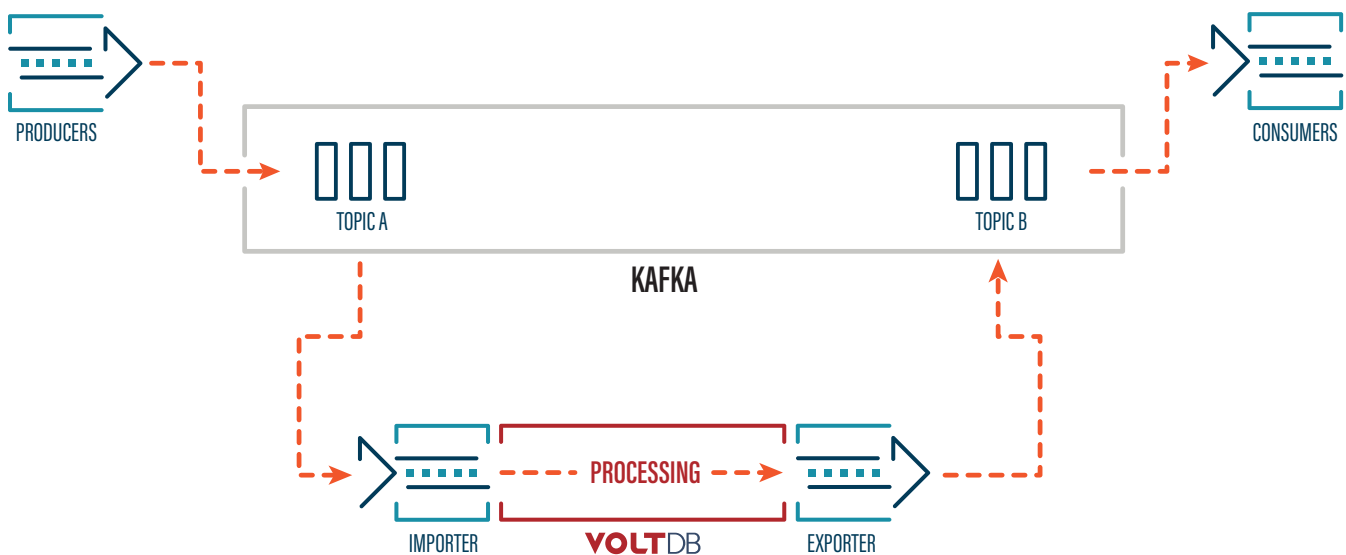
1. polling the external system to check for new data being available,
2. pulling any available data down,
3. ingesting this data event by event,
4. and then sending these events through the same stored procedures that applications would otherwise invoke directly in a traditional design.

Importers for different systems can be configured declaratively in a configuration file or through the VoltDB's user interface. These built-in importers start automatically whenever the database starts and stop when the database stops, making import an integral part of the database process. In addition, import connectors also provide durability to the events being ingested to ensure that no data is lost during disaster scenarios. Importer instances can be created on a running database for a no downtime integration.

The VoltDB Kafka Importer uses Kafka Consumer API to ingest data from multiple Kafka brokers and multiple topics with high performance. Developers can configure the importer using the same familiar properties as for a Kafka Consumer.

Export

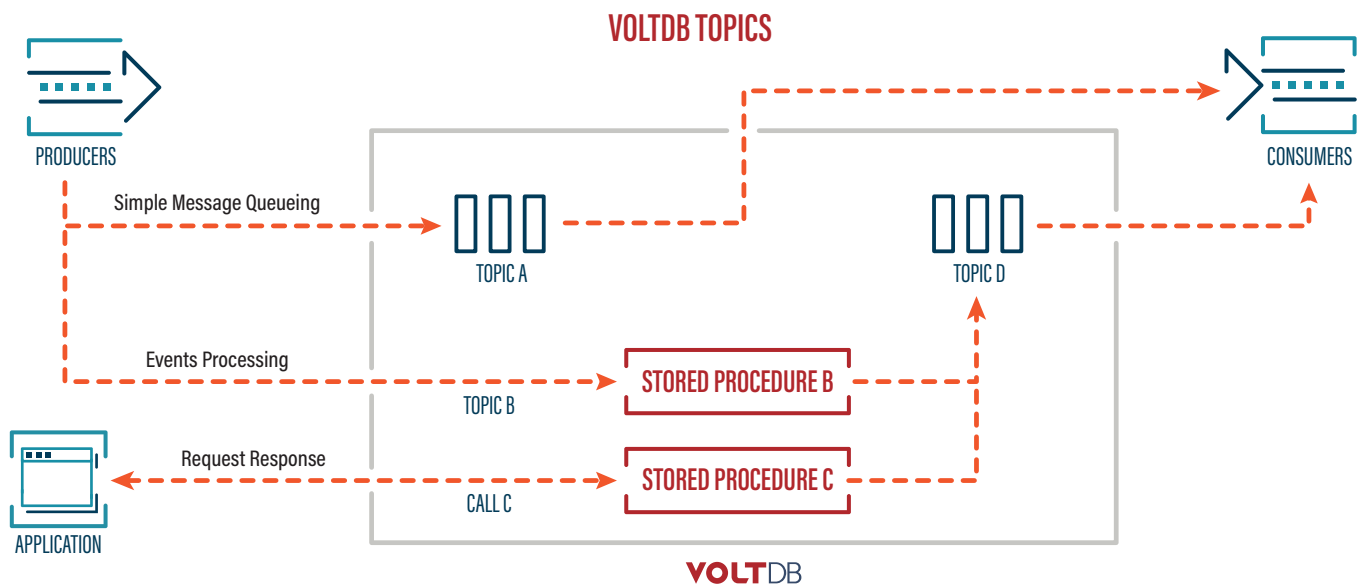
Export automates the reverse process from import, catching any data written to an export table or stream and sending it to the associated external target, whether it be a file, a service such as Kafka, or another database. Developers can choose to export specific records or to migrate rows that are being deleted from the tables due to expiration of TTL. The export process is transactional, so developers can be sure that no records will be lost during export. VoltDB guarantees that records will be exported at least once.



Replacement

VOLTDB TOPICS FEATURE (BETA)

Released in v10.1, the Topics feature allows applications to use publish and subscribe semantics to connect with VoltDB. In addition to providing better semantics for integrating with other systems, this feature also allows VoltDB to be used as a messaging queue—particularly useful when coupled with the strong data processing core that VoltDB was already built for. By standardizing on the Kafka consumer API and producer APIs to expose this feature, VoltDB can be easily dropped in as a replacement for Kafka especially for data streams that require low latency transaction processing.

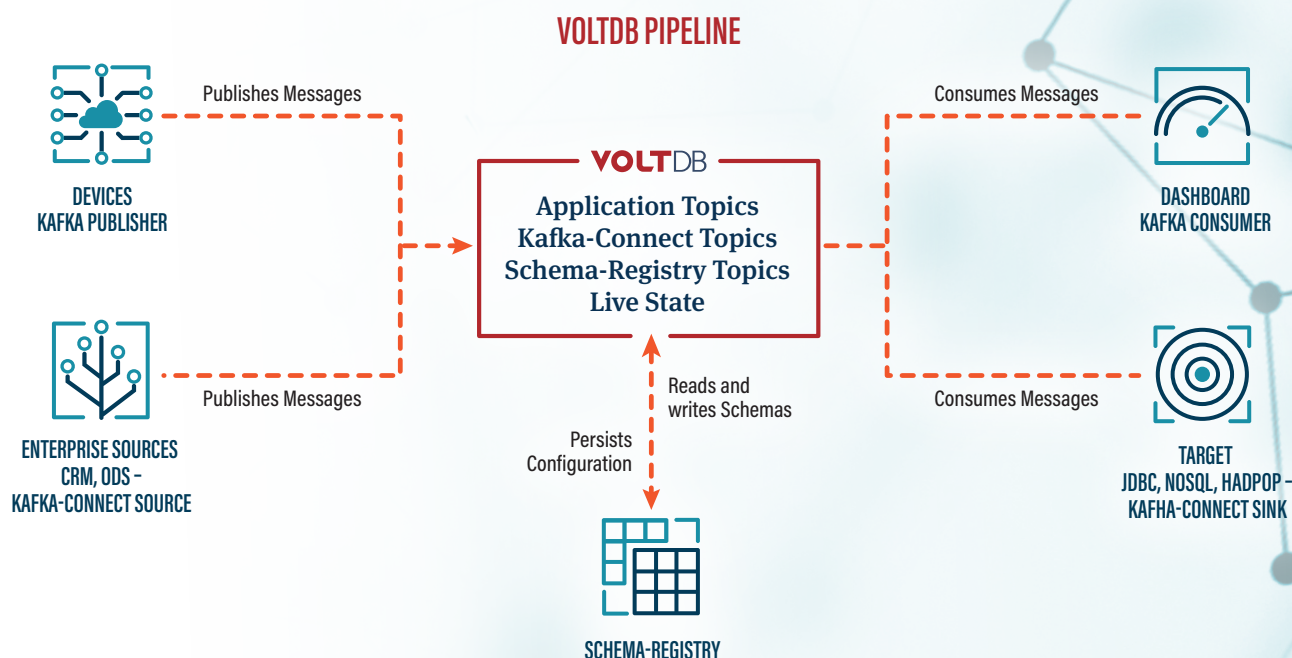


The benefits this delivers to applications are:

- Lower infrastructure costs by reducing the size of the Kafka cluster
- Better end-to-end latency due to lesser network hops
- A simplified architecture
- Ability to leverage Kafka Connect ecosystem

One advantage of VoltDB's Topics implementing the Kafka API is the ability to use Kafka Connect, which provides integration of other popular technologies with Kafka. Application deployments that look to have fewer moving parts in their architecture can embed publishers and consumers into their systems directly, while deployments that desire easy integration can take advantage of the rich Kafka Connect ecosystem.

Developers can easily wire together their data pipelines by using Source and Sink Kafka Connectors to integrate various systems together with VoltDB to take the unique advantage of having a message queue combined with a complex decisioning engine.



CONCLUSION

The worlds of big data and 5G are colliding with legacy infrastructures, to the point where even those that weren't considered "legacy" are quickly becoming legacies because they can't handle what's being thrown at them. Built in 2011, Kafka was born in the pre-5G, pre-IoT age, where applications could get by with a certain level of querying capability.

But now—things have changed. Kafka is still a powerful continual processor of streaming data, but to work well for real-time, in-event decisions based on complex streaming data, it requires some extra assistance.

VoltDB makes it easy to integrate with Kafka however you are using Kafka. Since we believe in and understand the power of Kafka and the needs of our customers, we're always looking for more ways to integrate with key technologies like Kafka, which is why we created the Kafka importer and exporter tool and also why we have developed the new pub-sub feature.

The bottom line: whether or not you use Kafka, we want your database architecture to support and accelerate your business and its goals, not interfere with them. That's why we always provide options, and why we are always available to chat about your data architecture needs.

To talk with VoltDB today [➤ CLICK HERE](#)

ABOUT VoltDB

VoltDB powers applications that require real-time intelligent decisions on streaming data for a connected world, without compromising on ACID requirements. No other database can fuel applications that require a combination of speed, scale, volume and accuracy. Architected by the 2014 A.M. Turing Award winner, Dr. Mike Stonebraker, VoltDB is a ground-up redesign of the relational database for today's growing real-time operations and machine learning challenges. Dr. Stonebraker has conducted research on database technologies for more than 40 years, leading to numerous innovations in fast data, streaming data and in-memory databases. With VoltDB, he realized the full potential of tapping streaming data with in-memory transactional database technology that can handle data's speed and volume while delivering real-time analytics and decision making. VoltDB is a trusted name in the industry already validated by leading organizations like: Nokia, Financial Times, Mitsubishi Electric, HPE, Barclays, Huawei, and more.