

# SHIFTX

Business process improvement,  
powered by FaunaDB

## ABOUT SHIFTX

ShiftX is a process modeling and improvement product, optimizing ecommerce and other business processes via a visual modeling tool. It offers a SaaS delivered environment for describing processes in a visual language to find problems and bottlenecks. ShiftX customers run ecommerce, fulfillment, manufacturing, and other process-heavy businesses.




Figure: Simplified ShiftX process example.

## PROJECT OVERVIEW

ShiftX customers model their business process with an expressive visual tool. The model can be used to direct improvements and gauge effectiveness. Business processes steps can include multiple types, for instance: phone or email interaction with a customer, payment processing, and shipping alerts. The model can direct attention to bottlenecks and common errors, helping the business improve where it matters most.

The ShiftX team leans on their database to make complex data easy to work with. By selecting a database with an expressive query language and deep object modeling

capabilities, including object references and temporal storage, they are able to write queries to provide insights to their customers. Finding a relationship where for instance, delays in one step introduce errors in a later step, takes a database which understands the data.

## REQUIREMENTS & CHALLENGES

ShiftX's primary deployment is in the cloud, but they know larger customers will require on-premise installations due to policy constraints. For this reason their database must be available both in a managed cloud, and for on-premise installation. A cloud-only solution wouldn't work for the biggest deals, so ShiftX needed a solution that's easy to package with on-premise installs.

Process modeling is inherently complex, with a variety of steps, each with many different stages, and types of inputs and outputs. ShiftX needed relational capabilities in a database with native references between objects. When your database understands your data, queries and interactions are simplified, and native object references were key to their data model. They also appreciated the flexibility of NoSQL documents to represent the variety of data their customers offer.

A database that supports relationships among entities matters, but the legacy solutions don't meet the operational criteria, as ShiftX wanted a cloud-native database, with on-premise availability.

Most graph and relational databases are designed and architected as scale-up solutions, rather than as cloud-native horizontally scalable databases. And existing NoSQL solutions have an operational burden that is more than ShiftX wants to carry.

## WHY FAUNA

ShiftX chose FaunaDB and GraphQL for their process modeling and improvement product, which is delivered to their customers as a service.

FaunaDB's data model includes ACID transactions, consistent indexes, joins, and native object references, all stored with configurable temporal snapshot retention. ShiftX depends on FaunaDB's rich query capabilities to optimize database traffic. ShiftX GraphQL resolvers combine FaunaDB query fragments so that application servers communicate efficiently with the database in a minimal number of queries, an additional benefit to FaunaDB's composable query language. This allows them to make intuitive use of native object references while building out user interface components.

As a SaaS vendor with a growing user base, ShiftX needs to provision new resources dynamically. They are able to utilize FaunaDB's hierarchical multi-tenancy to add new customers on demand. Customer workloads are safely isolated, and provisioning new tenants requires only a database query. Metering individual tenants allows FaunaDB Cloud to itemize costs based on customer usage.

While many SaaS customer are comfortable with cloud deployments, some valuable enterprises will require on-premise deployments. Fauna Cloud is fully managed so ShiftX can scale their SaaS customer base on-demand. When a new customer signs up with ShiftX, an isolated FaunaDB database is created to manage their data, alongside other parts of the ShiftX stack. Because FaunaDB is also available on-premise, Shiftx is also able to sell to customers that require on-premise deployments, without changing any code or introducing operational challenges.

## RESULTS

By choosing a FaunaDB and GraphQL stack, ShiftX accelerates time to market. They are able to build easily today on a platform that grows with their requirements. They have the flexibility to run in the cloud, or deploy on premise.

FaunaDB's expressive query language and ACID transactions minimized our time to market, while FaunaDB's data temporality will simplify our development

- Eigil Sagafos, CTO & Co-founder, ShiftX