

MODERNIZE YOUR APPLICATIONS:

## How Fauna's Document-Relational Database Surpasses MongoDB

Application frameworks, deployment methods and use cases are constantly evolving. Fauna was created specifically to enhance the scale, performance, security and agility that modern applications require.

### The Challenge: Legacy architecture & restrictive data model

MongoDB Atlas falls short when supporting modern application infrastructures by requiring application teams to be “database aware.” Despite Atlas automating deployment, it still requires effort to size & configure, limits data modeling options via lacking relational capabilities, only offers legacy RBAC access control and becomes a source of performance risk and cost inefficiencies as applications scale. It’s time for a new way.

### The Solution: A database built to enable application development velocity

Organizations choose Fauna over MongoDB because it combines the relational qualities of ACID transactions, joins, and schema with the flexibility of a document database, delivered as a fully serverless cloud API, ensuring seamless scalability and zero operational overhead. Fauna’s Jepsen-validated was designed from the ground up to deliver a new experience for application teams by offering an out-of-the-box serverless database ready for production scale, flexible data models, advanced security and effortless infrastructure. With Fauna, teams can focus on developing value-added features.



#### Deployment

##### Serverless-First with Zero Ops.

Eliminate tuning, provisioning, configuring, & capacity planning. Fauna is delivered as a Cloud API with auto-routing and active/active replicas it removes connection pools or locality / targeting and from consideration. With a service-mesh and true 2-tier model Fauna scales horizontally from zero to peak without sharding or partitioning.



#### Security

##### True Principle of Least Privilege with ABAC.

Upgrade from MongoDB’s collection-level RBAC. Secure data at its source with Fauna’s advanced ABAC system where both data and requestor’s attributes are available during both role grant logic as well as functions.

Fauna implements databases as containers and offers parent/child database relationships for simplified multi-tenancy and database-level isolation.



#### Performance / Scale

##### Maintain Low Latency at Internet Scale.

Avoid slowing down as you scale or needing to shard. Without the need to pre-provision clusters or configure per use case, Fauna can deliver low latency operations up to huge scale without the need for sharding. With its Calvin-based distributed transaction engine it can maintain low latencies over 3 cloud regions.



#### Data Modeling

##### Support Document and Relational Access Patterns.


Simplify application logic with native relational support. Fully enable normalized data models with Fauna’s dynamic query joins, schema enforcement and native ACID compliance for multi-document operations.

Quickly iterate features and move from dev to prod with schemaless but lock down with advanced schema enforcement paired and online operations.




Customer Success


Since Fauna documents do not have to follow a specific schema, we are able to save time and effort -- but we still can leverage Fauna's relational capabilities. It's the best of both worlds.

**Artur Tomusiak,**  
Senior Software Engineer 


Fauna and Fastly have been instrumental in our ability to test and release features quickly. This stack also enables us to deliver ultra-low latency because it's fully distributed and API-driven.

**Isis Baulig, CTO** 

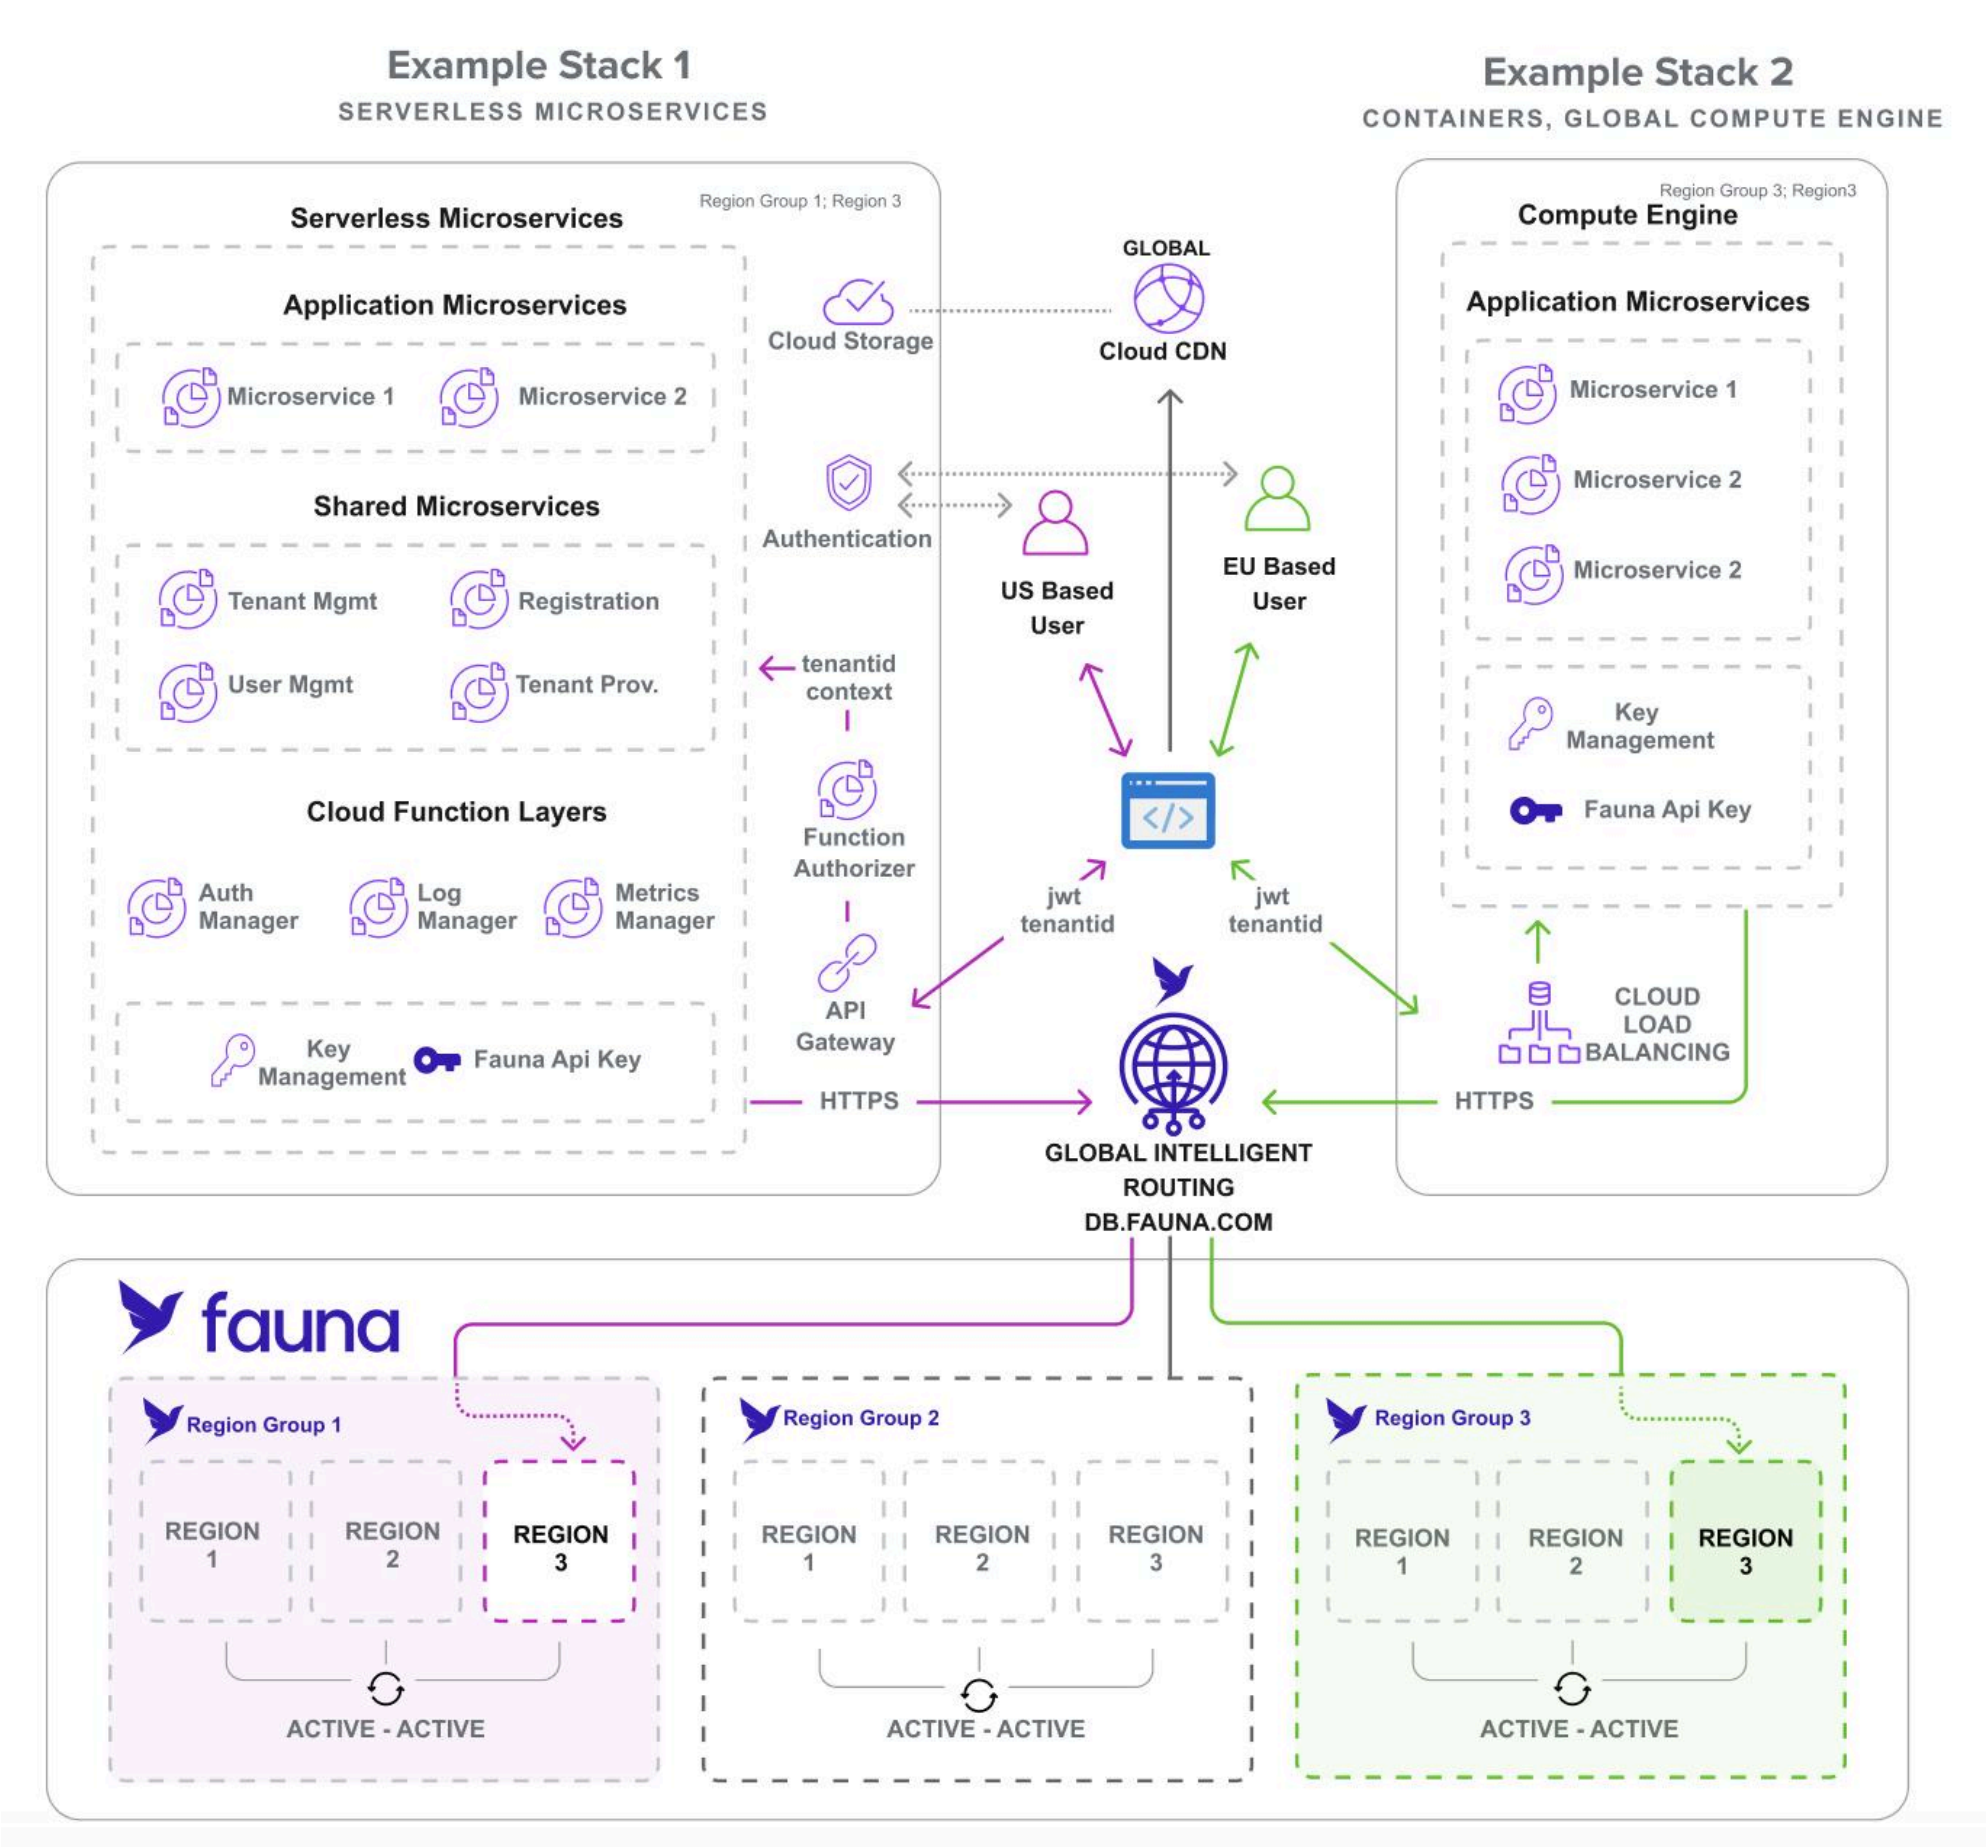
We needed a database that could support a distributed, multi-tenant architecture with robust ABAC and user controls. We looked at Dynamo and Mongo, but only Fauna delivered it all without extensive engineering.

**Arjun Bhatnagar, CEO** 

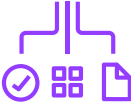
All of our web-based properties are built on Fauna and Cloudflare Workers, because they are both distributed-by-default and delivered as APIs, so they scale as we grow.


**Cameron Bell, Software Manager** 


Modern Architecture





Fully Featured


- 

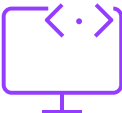
**Document-Relational Data Model**  
Combines document flexibility with relational querying, schema enforcement, and ACID transactions.
- 


**Multi-Active Serverless Engine**  
Auto-distribution, replication, multi-active reads & writes, & strong consistency. Eliminate server & cluster management, capacity planning, memory allocation, sharding, & more.
- 

**Cloud API Connectivity**  
Secure, modern HTTPS delivery with intelligent routing. Stateless, lightweight, and performant.
- 

**Schema as Code**  
Start schemaless and add structure that can be enforced over time with zero-downtime migrations.
- 

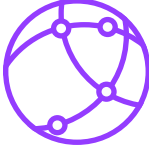
**Modern Security Model**  
Stateless, token-based authentication and dynamic ABAC offers identity and context-based access for more fine-grained and secure control
- 

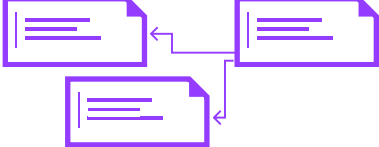
**Native Database Multi-Tenancy**  
Provide a secure container for each business or consumer your application serves through a native parent-child database model
- 

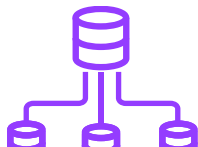
**Real-time Event Streaming**  
Enable live state management and power secure, modern and scalable interactive user experiences without the overhead of polling
- 


**Industry Compliant**  
SOC2 Type 2, GDPR, PCI, PII, & HIPAA certified.


Common Use Cases

- 

Multi-Region or Global Workloads
- 

Dynamic + Relational Access Patterns
- 

Multi-Tenant Apps
- 

Distributed Edge Workloads
- 

Serverless Architectures