



WELCOME

MongoDB Atlas

Much more than just a database



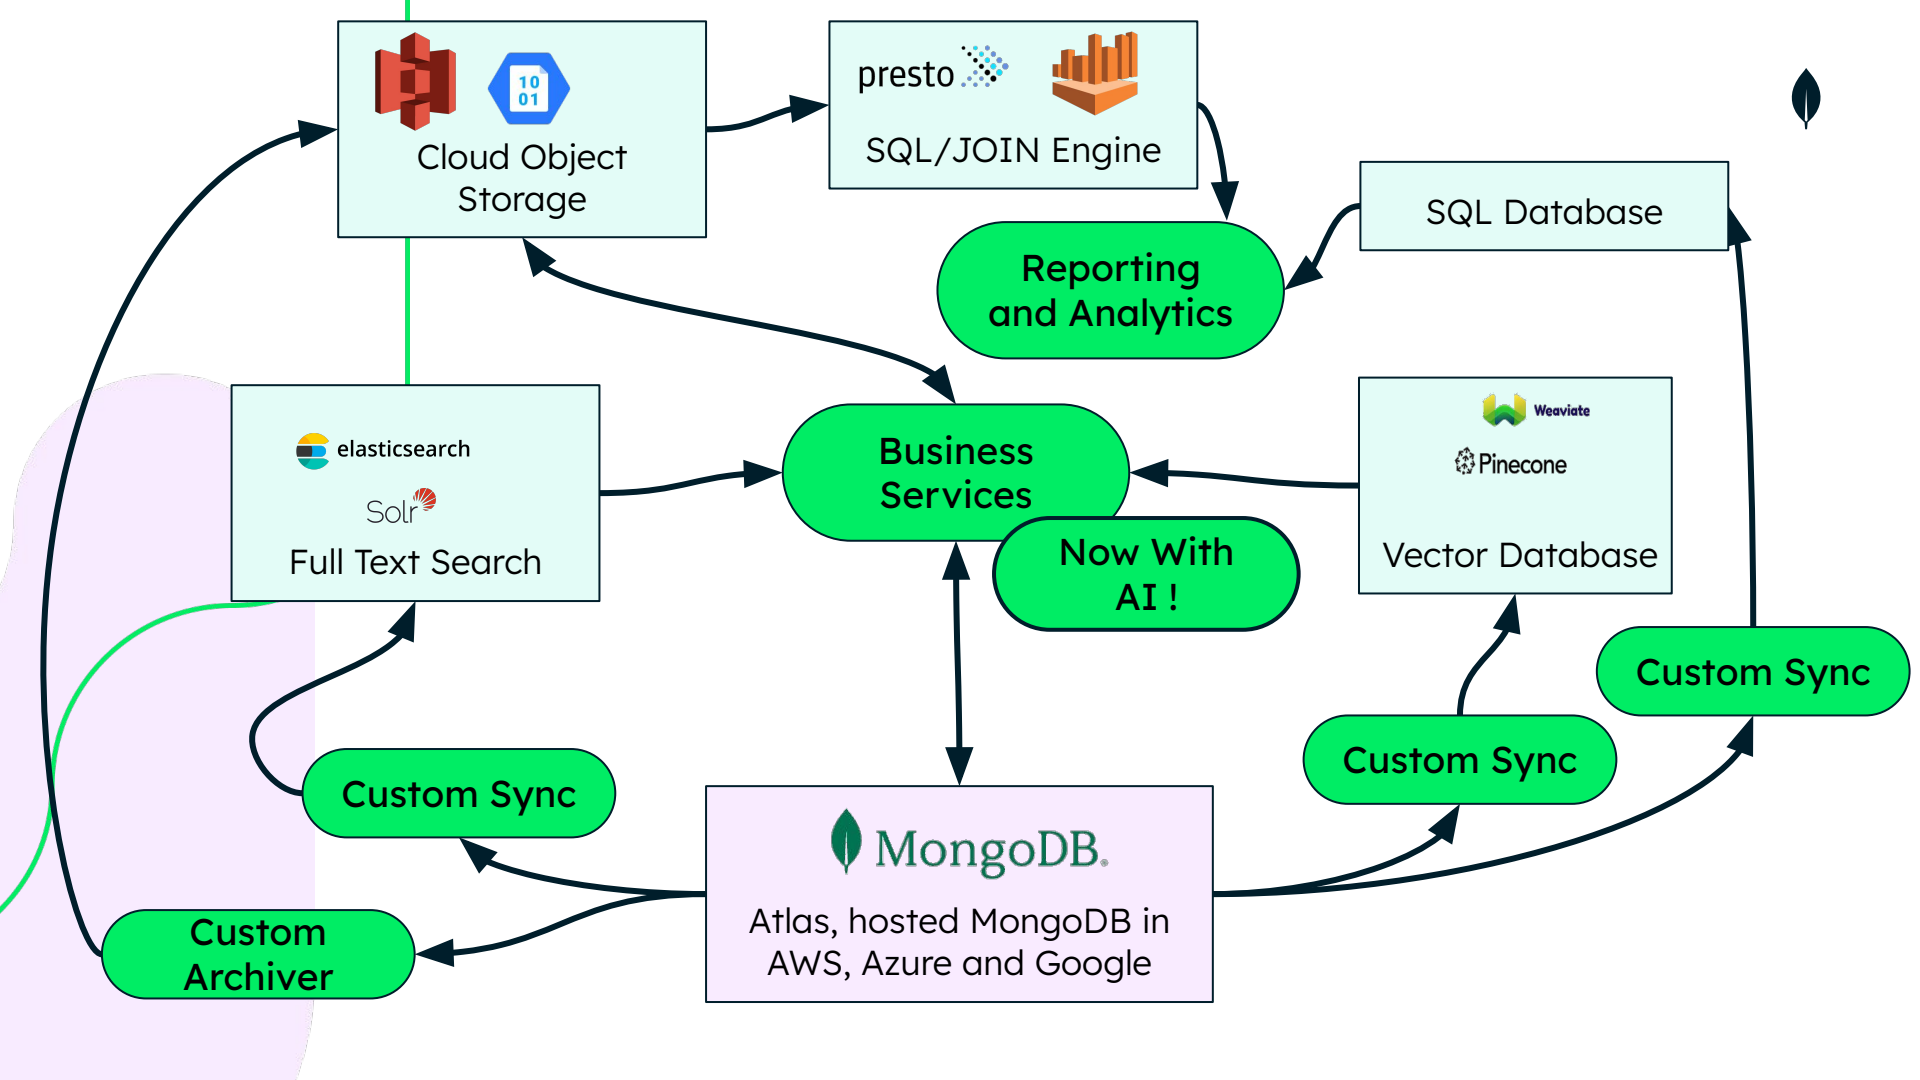
Andrew Morgan
Senior Staff Developer Advocate
clusterdb.com




MongoDB Atlas

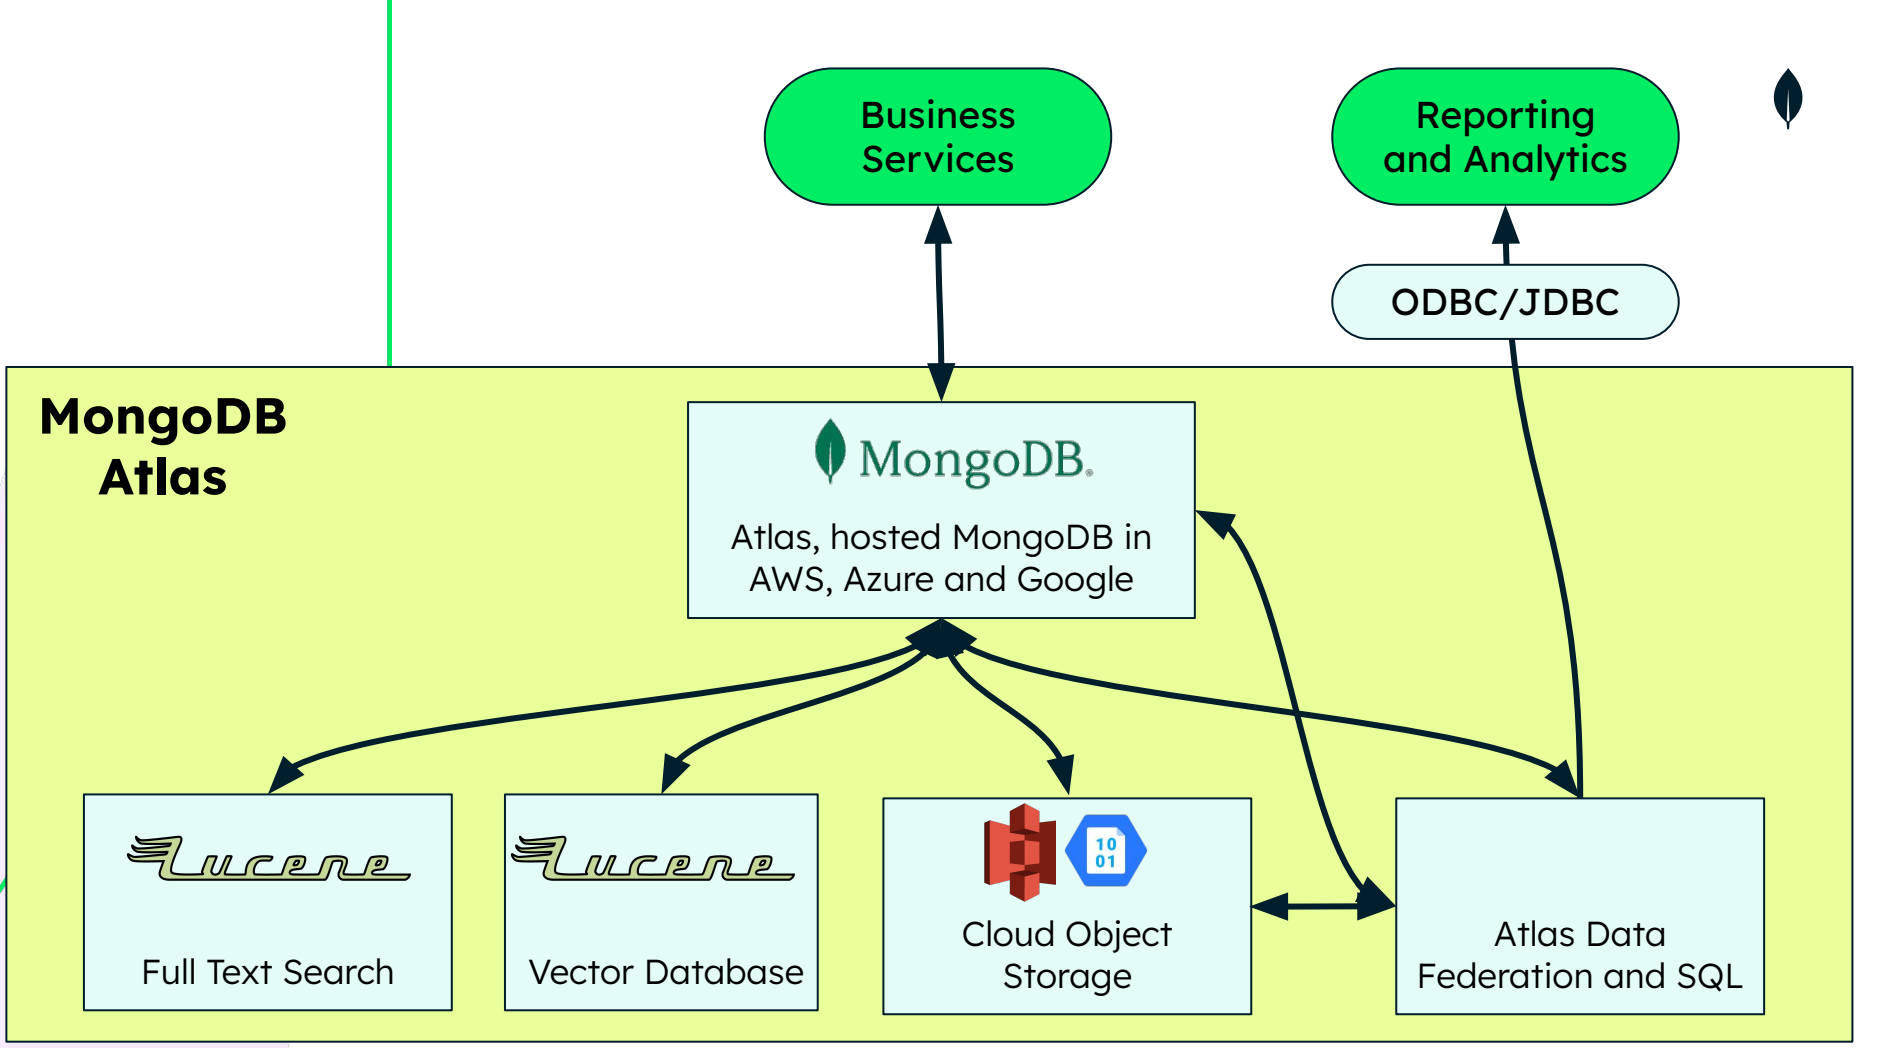

Developer Data Platform

"One Platform, One Connection, One API,
Many Data Services"





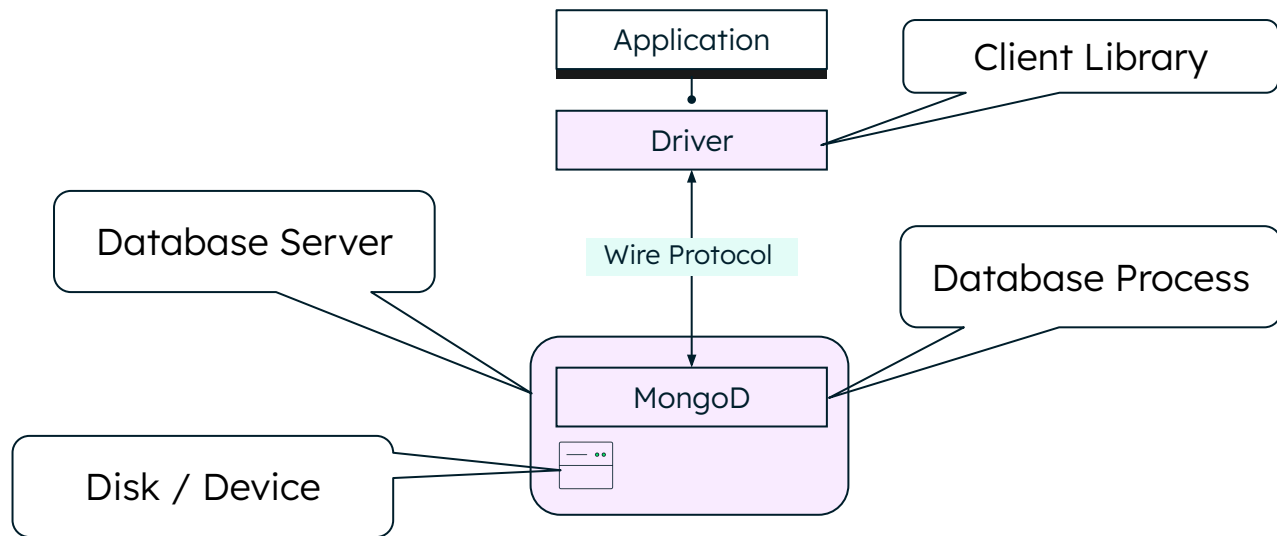
MongoDB Atlas



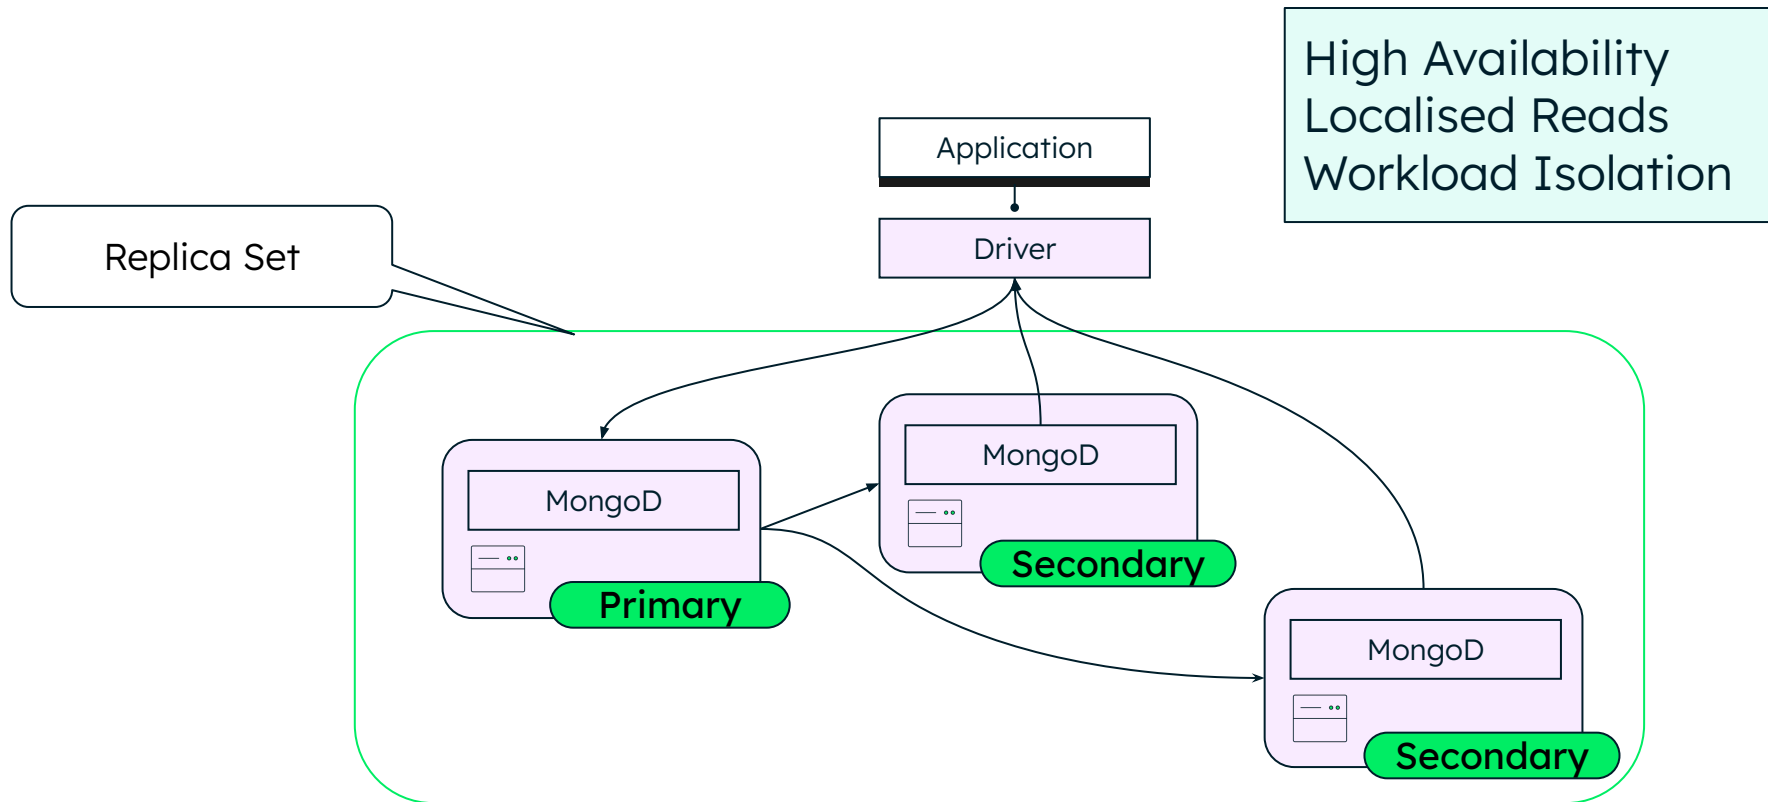


MongoDB Atlas DBaaS

MongoDB architecture



Replica Sets




Sharded Clusters

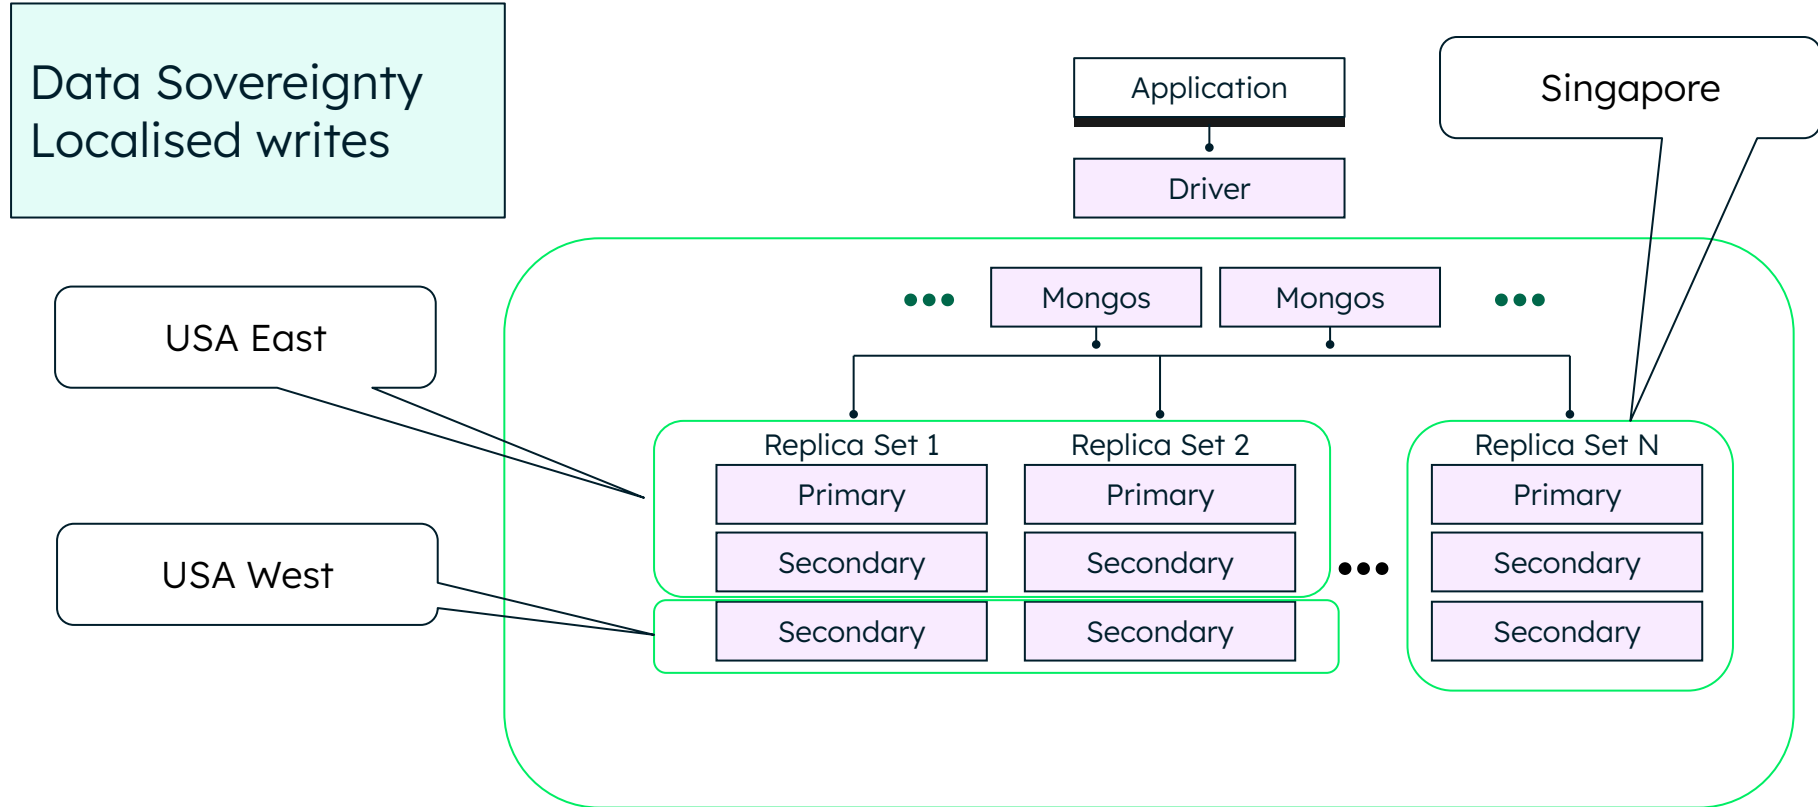



Sharded Cluster

Horizontal Scaling



Geo Sharded Clusters






Fine-grained monitoring & alerting

Create dashboards from dozens of throughput, performance, and utilization metrics to analyze historical database performance down to one minute granularity

Atlas projects come with a set of pre-configured alerts to cover common performance issues

Add and customize notifications when operations, server usage, or other conditions reach defined thresholds to suit your workloads

Get alerts over email, text, and other channels you already use



The Real-Time Performance Panel provides unparalleled insight into your database by allowing you to see cluster metrics in real-time, including:

- Operations
- Read/Writes
- Network In/Out
- Memory
- Hottest collections
- Slowest operations

Real-time performance metrics



Diagnose and resolve issues quickly with the help of the Query Profiler, which surfaces long-running queries and their key performance statistics

Review historical operations to quickly identify outliers and query patterns across collections and operation types

Always-on and available with up to 24-hour visibility into past operations

No need to download, process, or analyze server logs

Query Profiler


Your Infrastructure

Java Application


Node Application

Python Application

Atlas



MongoDB Cluster



MongoDB Cluster

Atlas Forms

Vacation Properties

Search Clear Edit Save Logout

Name	Beds	Address Country	Address Market
Copacabana Apartment Posto 6	3	Brazil	Rio De Janeiro
Cozy house at Beyoglu	1	Turkey	Istanbul
Room w private toilet 3 min to MTR!	1	Hong Kong	Hong Kong
Quarto de Casal em Ipanema	1	Brazil	Rio De Janeiro

Id: 10038496 Name: Copacabana Apartment Posto 6 Property Type: Apartment Room Type: Entire home/apt

Bed Type: Real Bed Minimum Nights: 3 Maximum Nights: 75 Cancellation Policy: strict_14_with_grace_p Last Scraped: 11/02/2019, 05:00

Calendar Last Scraped: 11/02/2019, 05:00 First Review: 18/01/2016, 05:00 Last Review: 28/01/2019, 05:00 Accommodates: 4 Bedrooms: 1

Beds: 3 Number Of Reviews: 70 Bathrooms: 2.0 Amenities: TV Price: 119.00

Security in Public Cloud






Secure wherever you put your data

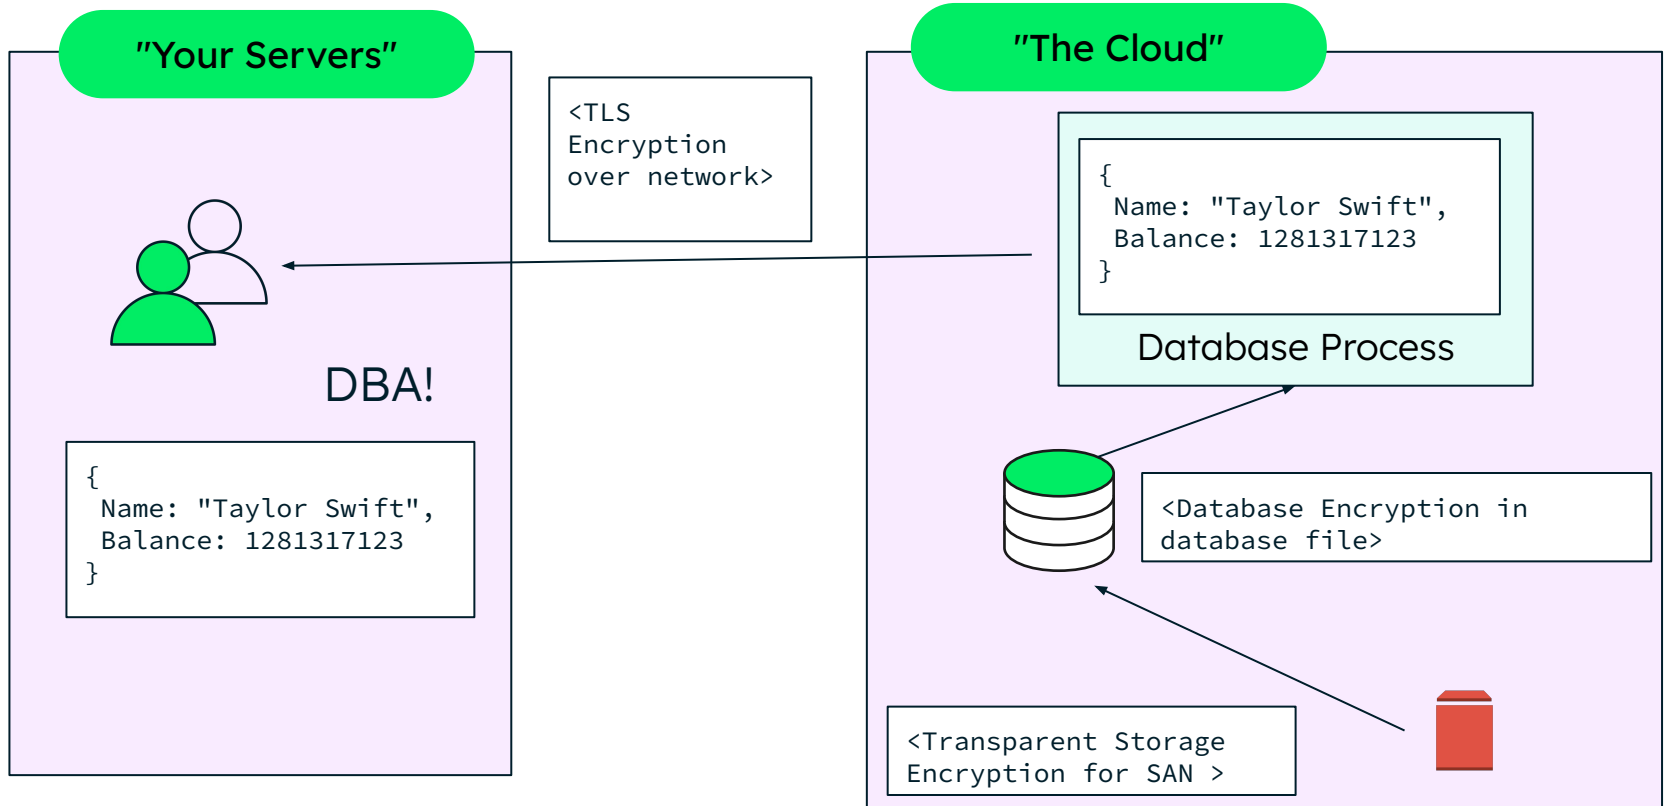
All MongoDB Atlas customer projects are deployed into their own VPC for network isolation. Private network peering is available for databases on all three clouds.

Business Trust Needs	Security Features
Authentication	SCRAM, X.509, LDAPS, AWS IAM, IP Access Lists
Authorization	RBAC, Read-Only Views, Field-Level Redaction
Auditing	Admin, DML, DDL, DCL, Role-based
Encryption	In-Flight: TLS 1.2+ At-Rest: Encrypted Hardware, Volume and Database Storage Engine (AES-256) Key Management: Cloud Services (AWS KMS, Azure KV, Google Cloud KMS)

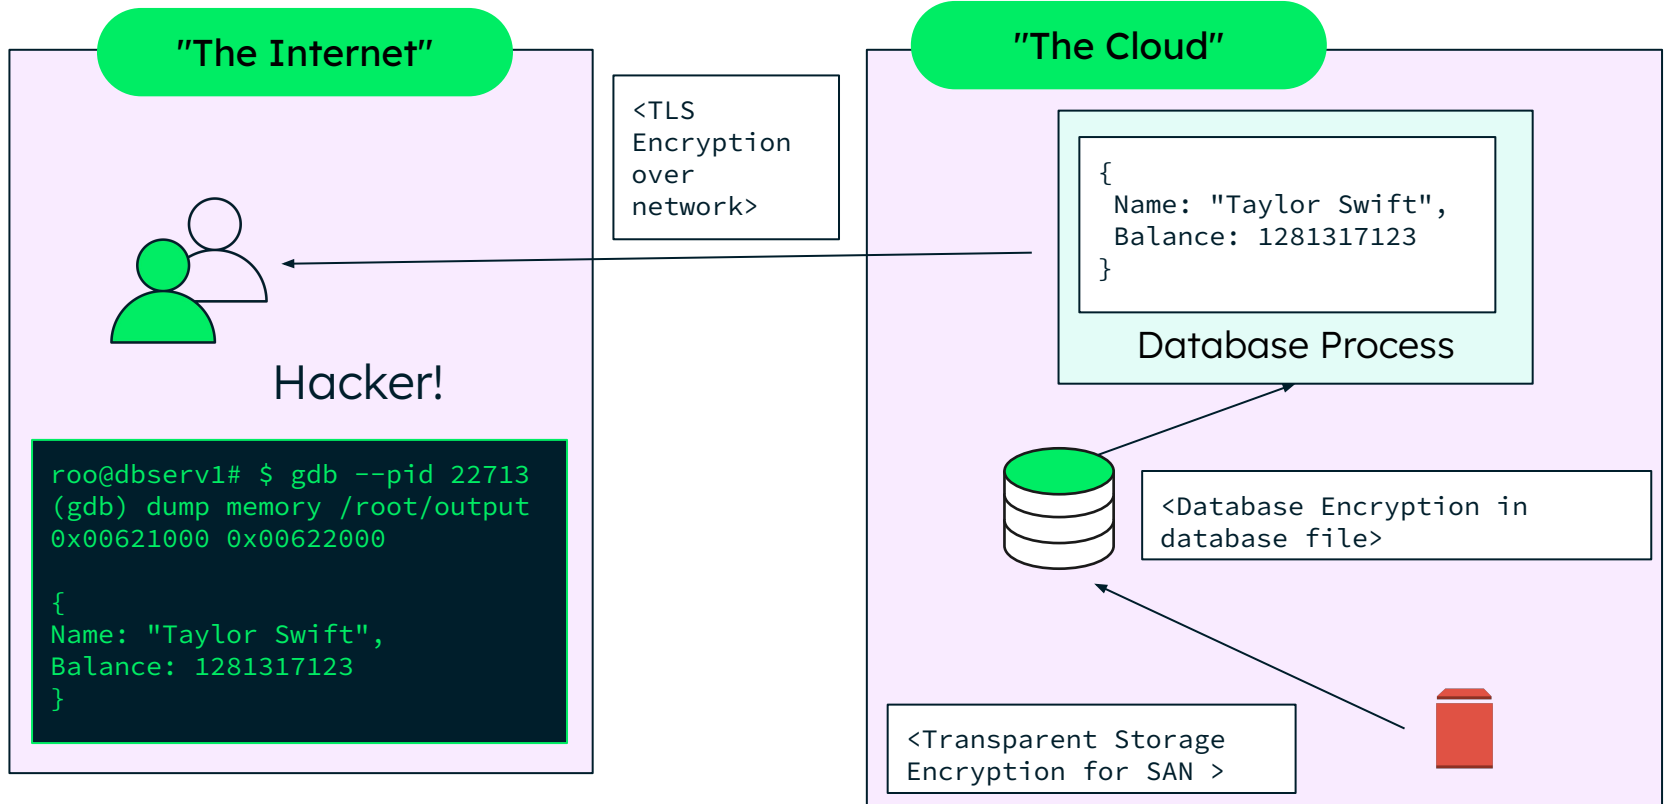
Encryption in Use



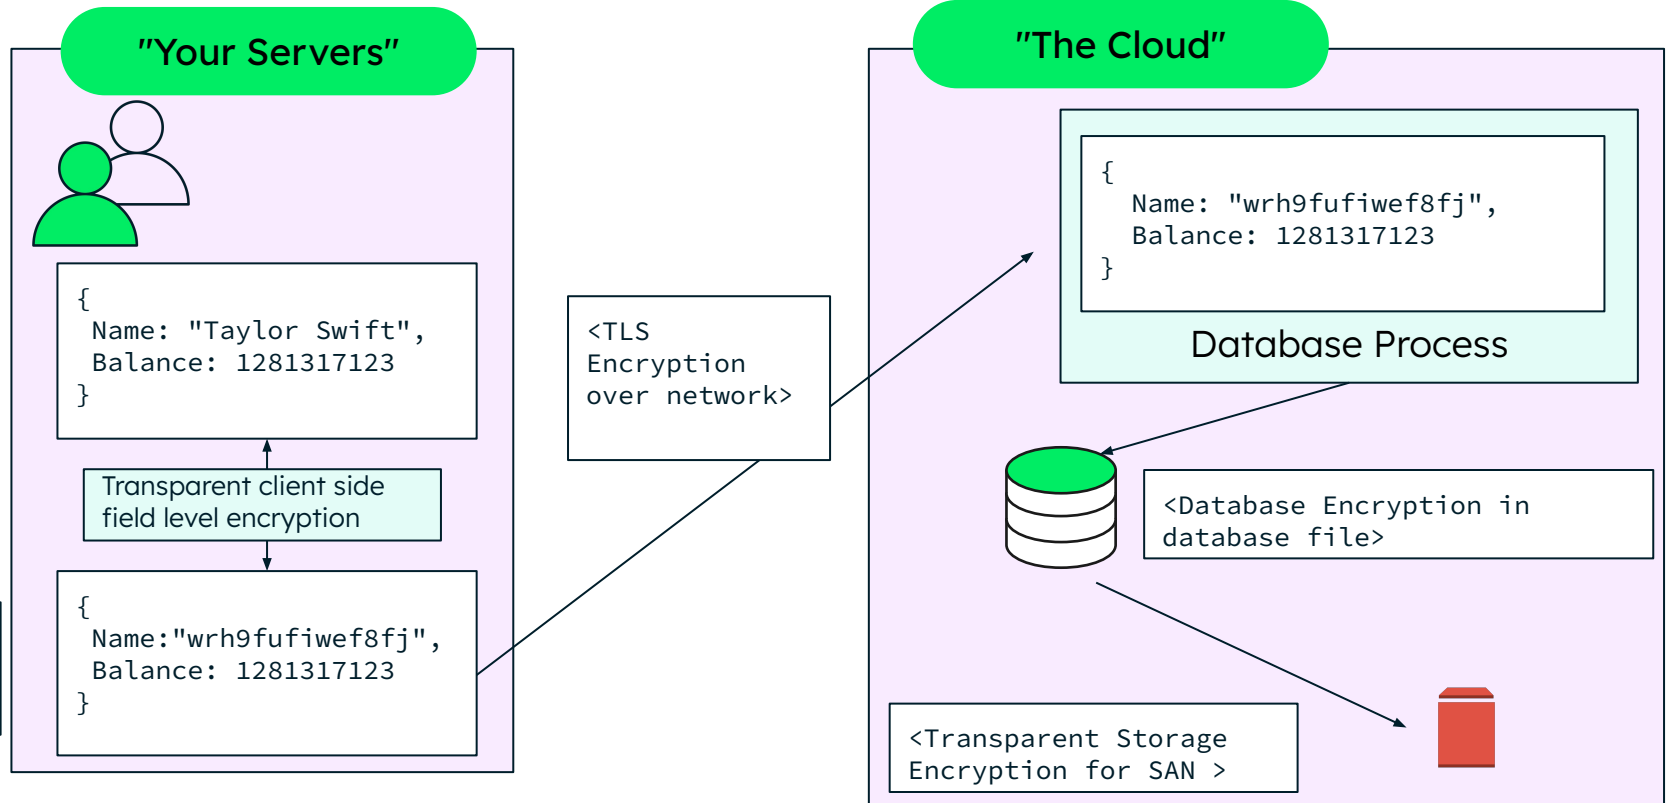
Encryption in Use - Threats



Encryption in Use - Threats

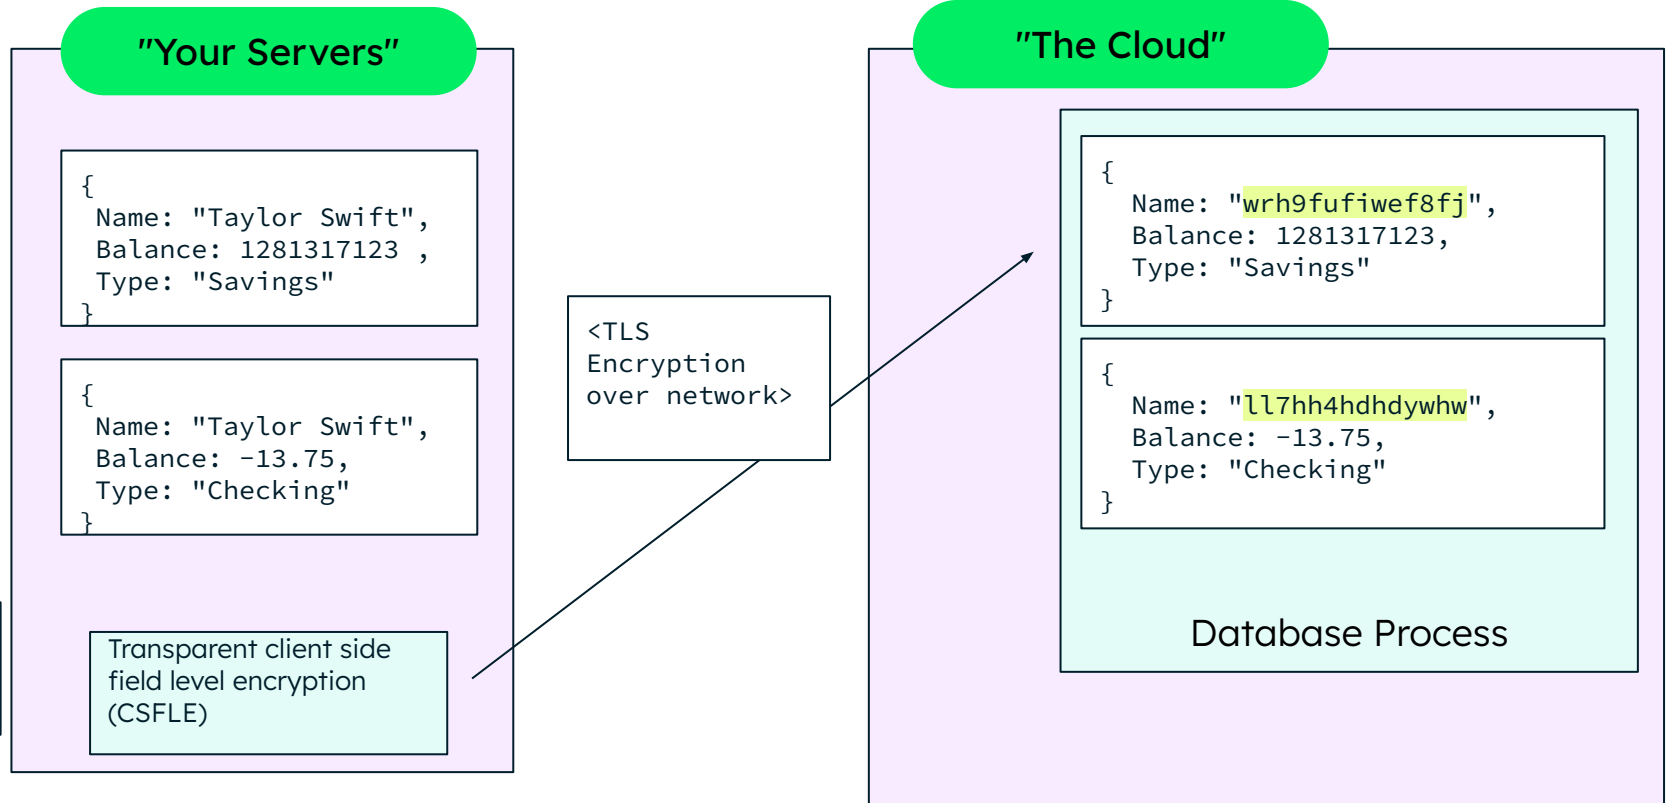


Encryption in Use - The Fix







CSFLE is Non Deterministic



Queryable CSFLE

“Any sufficiently advanced technology is indistinguishable from magic.”





Atlas Search


Embedded full-text search in MongoDB Atlas that gives developers a seamless and scalable experience for building relevance-based app features

- Leverage the power of **Apache Lucene**, the world's leading search engine technology
- Eliminates the need to run a separate search system alongside your database

How developers build app search today




Stand up a replication mechanism to keep the systems in sync



How developers build app search today



Architectural complexity in your application stack



✗ Developer friction

Different query APIs and drivers for database and search, coordinate schema changes


✗ Pay the sync tax

Requires its own systems and skills. Recovering sync errors can consume 10% of a developer's time

✗ Operational overhead

More to provision, secure, upgrade, patch, back up, monitor, scale, etc.


Rich Text search features



Atlas Search offers the features needed to fine-tune your search results to help users find what they need, including

- Fuzzy search
- Autocomplete
- Filters and facets
- Synonyms
- Multi-data type support
- Multi-language support
- Highlighting
- Rich query language


Query Relevance Ranking



Atlas Search allows you to weight each clause of a query as well as define the score for fuzzy matches to score and sort results

- Weights for OR clauses
- Weights for nearness to a value
- Weights for nearness to a geolocation
- Weights for how well a text search matches

Termlist Indexes with Intersection



Atlas Search allows you to efficiently query by any combination of N indexed values

- Sorted lists not BTrees for easy intersection
- Allows you to define N searchable attributes and allow any combination
- Can be Ranked or Boolean



Your Infrastructure

Java Application

Node Application

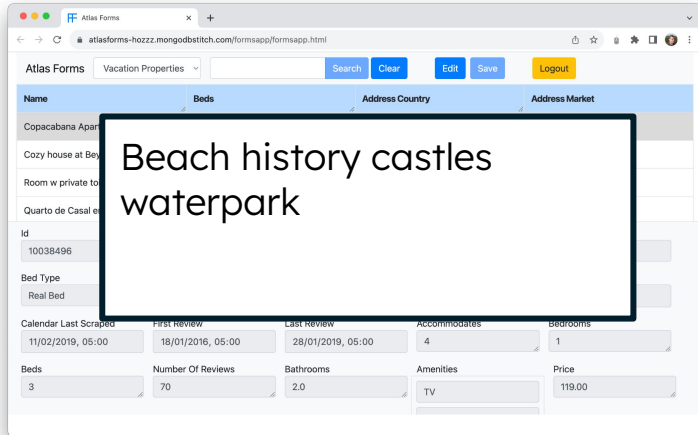
Python Application

MongoDB Cloud




Atlas Search Indexes

Lucene



Vector indexes and Search



Vector Search allows you efficiently compare vectors (arrays of doubles)


- Euclidian, Cosine or dotProduct
- Up to 4096 dimensions



What is a Vector?

RED

Red



Pantone®
485C

Crimson



"A glowing ball of
molten iron"

Merah

"Light at
625–740
nanometres
wavelength "

Rouge

Röt


R G B

[255, 0, 0]

C M Y

[0, 100, 100]

Generating Vectors



The accident involves two vehicles: a sedan in front and an SUV behind it. The front car has visible damage to its rear bumper, which is crumpled, while the SUV behind it has a damaged front end, with the hood slightly bent and the front crumpled.

Several people are exiting the sedan, appearing distressed, likely in response to the collision.

"Car Crash"



AI Model with
Vectorising Function:

- OpenAI, AWS Titan, Claude, Llama etc.
- Converts data to seemingly random but comparable vectors.



1500 element array
of numbers




Similar 1500
element array of
numbers



Similar 1500
element array of
numbers

What is Vector Search?

- Locating nearest neighbor to query vector out of a set of vectors
- “Nearness” determined by a distance metric, such as cosine similarity or euclidean distance
- Imagine comparing an array of colours easy to do in linear time hard to do efficiently!
- MongoDB uses an optimized approximate K Nearest Neighbour algorithm



[Source: Open Data Science](#)



Your Infrastructure

Java Application

Node Application

Python Application

MongoDB Cloud

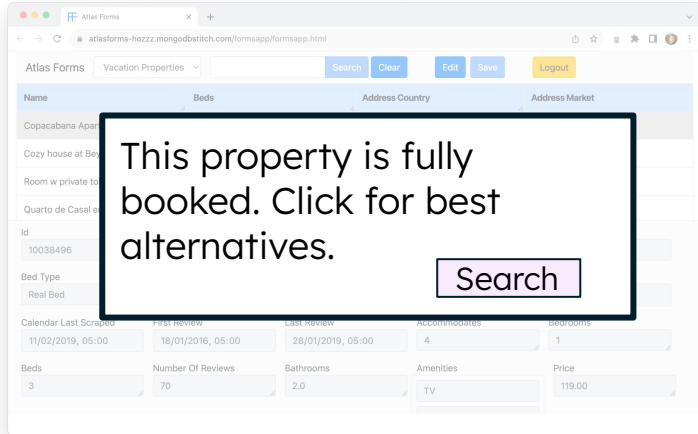


Atlas Vector Search Indexes

Lucene

This property is fully booked. Click for best alternatives.

Search



Retrieval Augmented Generation (RAG)

Generative AI will make things up if it doesn't know the answer (Hallucination) or give generic answers

Perform a "fuzzy" query against a private/internal database of facts using vector search.

Include those facts with your request to a generative AI to help it answer correctly based on non public knowledge.

Using OpenAI

I have two teenage children, my wife likes to relax by a pool and I like to photograph and watch wildlife, where should we go on vacation. Keep recommendation to 100 words max.

Consider Costa Rica for your vacation. The country offers stunning beaches and resorts with pools perfect for relaxation. For wildlife enthusiasts, Costa Rica is a paradise, with numerous national parks like Manuel Antonio and Corcovado where you can photograph diverse species, from monkeys to exotic birds. The Monteverde Cloud Forest is also a fantastic spot for wildlife watching. Your teenagers can enjoy thrilling activities like zip-lining and surfing, ensuring everyone has a great time. Costa Rica's mix of relaxation, adventure, and natural beauty makes it an ideal destination for your family.

Using OpenAI + RAG

I have two teenage children, my wife like to relax by a pool and I like to photograph and watch wildlife, where should we go on vacation. Keep recommendation to 100 words max.

I would recommend **Sabi Sabi Earth Lodge in South Africa** for your vacation. It combines your love for wildlife with the luxury of a serene pool where your wife can unwind. The lodge is situated in a private game reserve, offering incredible safaris where you can photograph the Big Five, while your teenage children can enjoy outdoor activities in a stunning natural setting. It's a perfect blend of adventure and relaxation, ensuring everyone in the family has an unforgettable experience.

Vectorize Query

0.9870816601	0.9212518477	0.7938539846	0.6773050674	0.314733618	0.02072906426
0.3652777941	0.4383044711	0.7494721752	0.2249677096	0.8603055981	0.04720800252
0.01054231581	0.1235446668	0.5270369584	0.2976291243	0.365019294	0.8050087506
0.5467109651	0.0989005923	0.7964032171	0.02597752136	0.7754301998	0.1295777033
0.0980835799	0.7202874228	0.05005313817	0.4573264472	0.494603545	0.8243313108
0.2135911819	0.7897491464	0.03998570485	0.6758764308	0.994015403	0.2896613869
0.552871699	0.5112042851	0.5502367792	0.9231493791	0.8606615835	0.6588229328
0.1944920471	0.7945369336	0.6520088137	0.4635946474	0.09162145717	0.2000617961
0.01888031182	0.2748397977	0.6829929143	0.9728749156	0.4403009302	0.6540594532
0.6341658435	0.2849768713	0.7811158334	0.8320092282	0.7106184085	0.5464651854
0.1984323363	0.9172822676	0.7725227953	0.5383951046	0.3963449666	0.1313129828
0.802917501	0.01069969709	0.1092702249	0.1183206045	0.058687351	0.8016996534
0.6914953664	0.9222376878	0.2134253353	0.7601409617	0.5654756848	0.2668963636
0.8485181796	0.05679064608	0.2110680153	0.4666821241	0.9637401599	0.114011322

Match internal database records

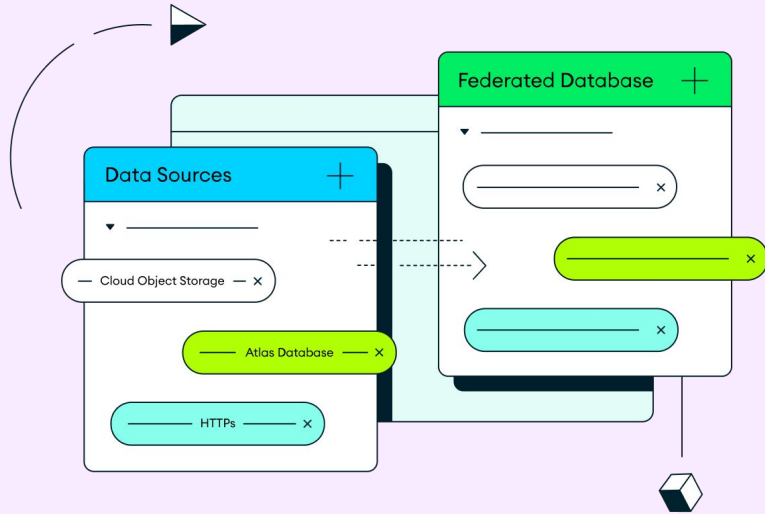
' **Anantara Golden Triangle Elephant Camp & Resort, Thailand** - Overlooking the lush jungle, this luxury resort offers close encounters with elephants, river tours, and a beautiful pool for relaxation.
Finca Rosa Blanca, Costa Rica - Nestled in a coffee plantation, this eco-friendly hotel offers birdwatching tours, nearby cloud forests, and a sunny pool with mountain views.
Sabi Sabi Earth Lodge, South Africa - Situated in a private game reserve, this lodge offers Big Five safaris, outdoor activities, and a serene pool with panoramic views of the bushveld.
The Brando, French Polynesia - Located on a private island, this resort features snorkeling with marine life, sun-kissed beaches, and a luxurious pool surrounded by tropical wildlife.
Inkaterra Reserva Amazonica, Peru - This eco-lodge in the Amazon rainforest offers guided wildlife tours, canoeing, and plenty of sunshine with riverside relaxation options. "

+ "please offer from" +



Atlas Data Federation Online Archive and Data Lake Pipelines

Atlas Data Federation

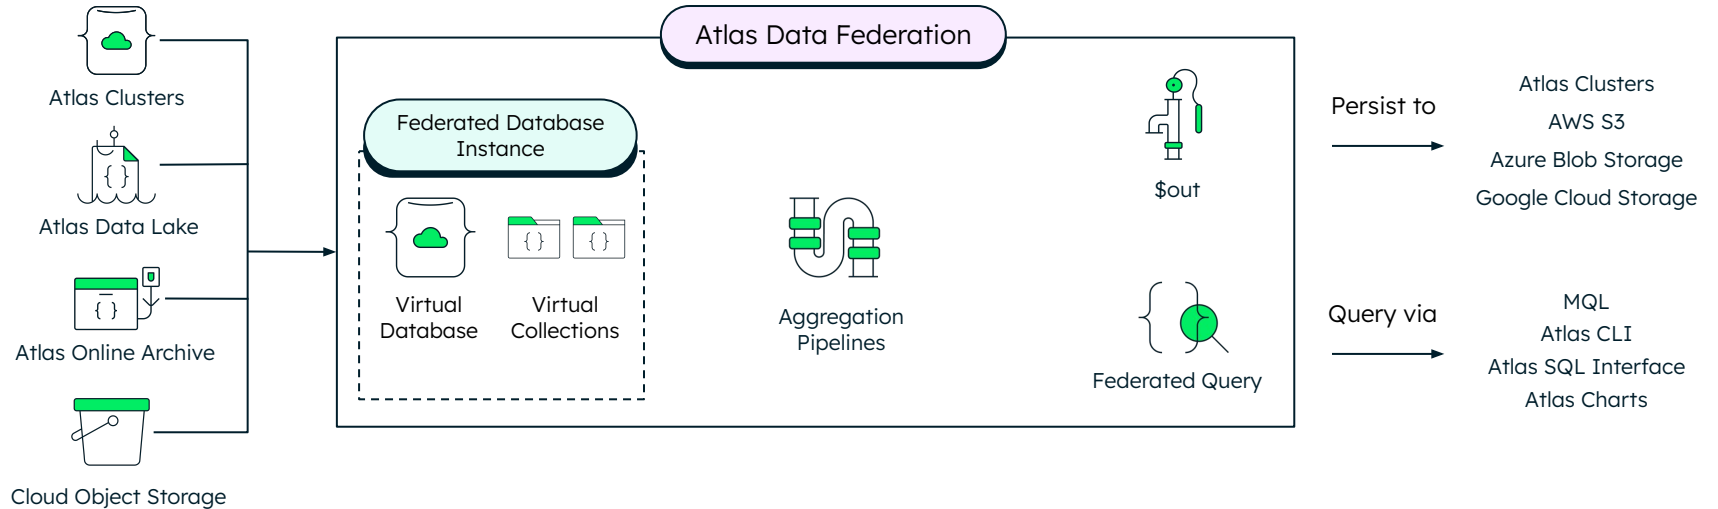


Data Federation gives Atlas users the ability to **combine, enrich, and transform data**, execute **federated queries** across data sources and **output** data to their preferred destination, through a stable, flexible, performant service.





How does Data Federation work?



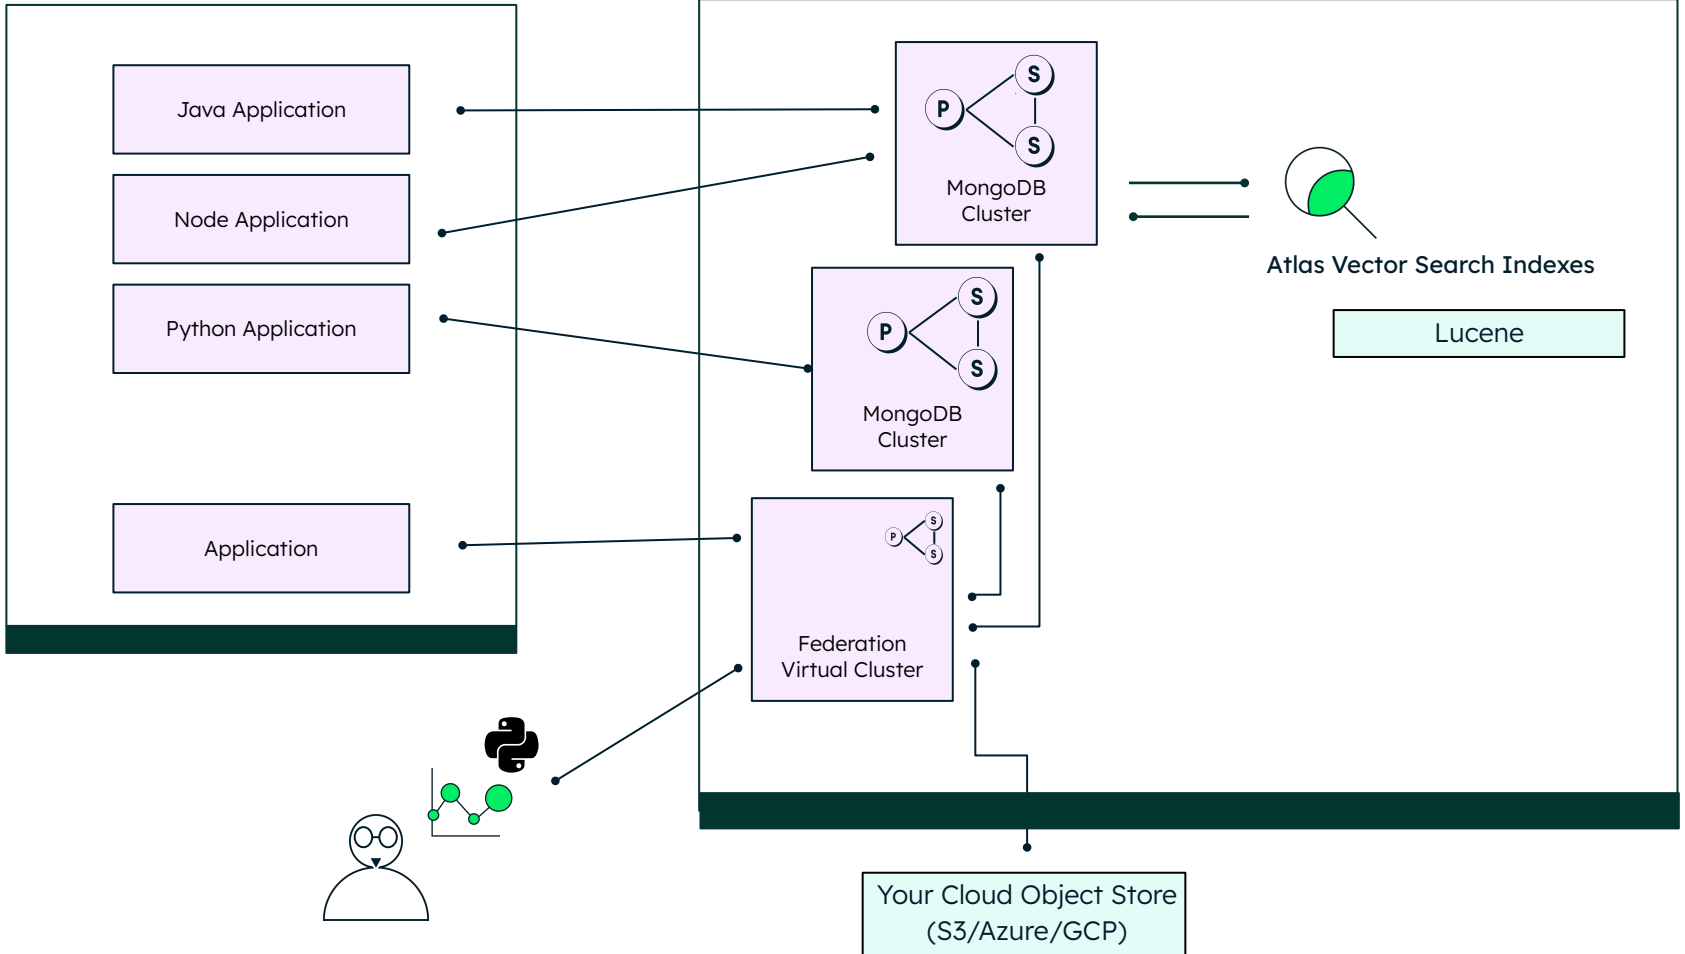
Configure a Federated Database Instance

Combine, enrich, and transform data

Persist data into a new location or directly query it

Your Infrastructure

MongoDB Cloud






Atlas SQL Interface, Connectors, and Drivers




Atlas SQL Interface, Connectors, and Drivers

Leverage SQL tools to easily query and generate rapid insights from MongoDB Atlas






Why Atlas SQL Interface




Easily Connect To Atlas

Natively built connectors provide easy setup and install with distinct guidance for Custom Connectors



First-class SQL experience

Built with a MongoDB-native, SQL-92 compatible dialect, to get the most out of your SQL tools





Directly query live data

Ensure consistent analysis on live operational data and eliminate the need for ETL or data duplication

Your Infrastructure

- Java Application
- Node Application
- Python Application
- Pricing Optimisation Engine (as a Service)



MongoDB Cloud



Atlas Vector Search Indexes

Lucene

Charts



Thank you for
your time.

