WELCOME

Investor Session: Product Update

MongoDB.local NYC 2024

Michael Gordon
COO & CFO

Sahir Azam
CPO

Andrew Davidson
SVP Product

Paul Done
Field CTO

Safe Harbor

This presentation and the accompanying oral presentation have been prepared by MongoDB, Inc. ("MongoDB" or the "company") for informational purposes only and not for any other purpose. Nothing contained in this presentation is, or should be construed as, a recommendation, promise or representation by the presenter or MongoDB or any officer, director, employee, agent or advisor of MongoDB. This presentation does not purport to be all-inclusive or to contain all of the information you may desire. Information provided in this presentation speaks only as of the date hereof, unless otherwise indicated.

This presentation contains "forward-looking statements" within the meaning of Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended, including but not limited to statements regarding our financial outlook, long-term financial targets, product development, business strategy and plans, market trends and market size and opportunities. These forward-looking statements include, but are not limited to, plans, objectives, expectations and intentions and other statements contained in this presentation that are not historical facts and statements, and may be identified by words such as "anticipate," "believe," "continue," "could," "estimate," "expect," "intend," "may," "plan," "project," "will," "would" or the negative or plural of these words or similar expressions or variations. These forward-looking statements reflect our current views about our plans, intentions, expectations, strategies and prospects, which are based on the information currently available to us and on assumptions we have made. Although we believe that our plans, intentions, expectations, strategies and prospects as reflected in or suggested by those forward-looking statements are reasonable, we can give no assurance that the plans, intentions, expectations or strategies will be attained or achieved. In particular, the development, release, and timing of any features or functionality described for MongoDB products remains at MongoDB's sole discretion. Product roadmaps do not represent a commitment, promise or legal obligation to deliver any material, code, or functionality and you should not rely on them to make your purchase decisions. Furthermore, actual results may differ materially from those described in the forward-looking statements and are subject to a variety of assumptions, uncertainties, risks and factors that are beyond our control including, without limitation: the effects of the ongoing military conflict between Russia and Ukraine and Israel and Hamas on our business and future operating results; economic downturns and/or the effects of rising interest rates, inflation and volatility in the global economy and financial markets on our business and future operating results; our potential failure to meet publicly announced guidance or other expectations about our business and future operating results; our limited operating history; our history of losses; failure of our platform to satisfy customer demands; the effects of increased competition; our investments in new products and our ability to introduce new features, services or enhancements; social, ethical and security issues relating to the use of new and evolving technologies, such as artificial intelligence, in our offerings or partnerships; our ability to effectively expand our sales and marketing organization; our ability to continue to build and maintain credibility with the developer community; our ability to add new customers or increase sales to our existing customers; our ability to maintain, protect, enforce and enhance our intellectual property; the growth and expansion of the market for database products and our ability to penetrate that market; our ability to integrate acquired businesses and technologies successfully or achieve the expected benefits of such acquisitions; our ability to maintain the security of our software and adequately address privacy concerns; our ability to manage our growth effectively and successfully recruit and retain additional highly-qualified personnel; and the price volatility of our common stock. These and other risks and uncertainties are more fully described in our filings with the Securities and Exchange Commission ("SEC"), including in the section entitled "Risk Factors" in our Annual Report on Form 10-K for the fiscal year ended January 31, 2024 and in other filings and reports we may file from time to time with the SEC.

This presentation includes market and industry data and forecasts that the company has derived from independent consultant reports, publicly available information, various industry publications, other published industry sources, and its internal data and estimates. Independent consultant reports, industry publications and other published industry sources generally indicate that the information contained therein was obtained from sources believed to be reliable. Although the company believes that these third-party sources are reliable, it does not guarantee the accuracy or completeness of this information, and the company has not independently verified this information. The company's internal data and estimates are based upon information obtained from trade and business organizations and other contacts in the markets in which the company operates and management's understanding of industry conditions. Although the company believes that such information is reliable, it has not had this information verified by any independent sources. In addition, the information contained in this presentation is as of the date hereof (except where otherwise indicated), and the company has no obligation to update such information, including in the event that such information becomes inaccurate or if estimates change. Subsequent materials may be provided by or on behalf of the company in its discretion and such information may supplement, modify or supersede the information in these materials. Neither the company, nor any of its respective affiliates, advisors or representatives shall have any liability whatsoever (in negligence or otherwise) for any loss or damage howsoever arising from any use of these materials or their contents or otherwise arising in connection with these materials.

Except as required by law, we undertake no obligation to update any forward-looking statements included in this presentation as a result of new information, future events, changes in expectations or otherwise. Nothing in this presentation is, and should not be construed as, an offer to sell or a solicitation of an offer to buy, any securities.

We are pursuing one of the
largest & fastest growing markets
in software

Data Management Software Market, $Bn

Source: IDC, Data Management Software Market, $Bn

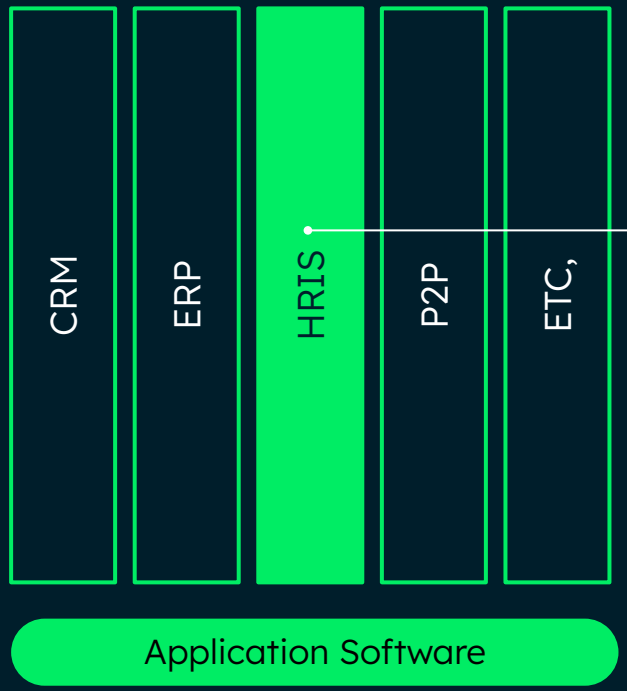

The unit of competition in many other software markets is a **customer**

Competition tends to be binary - only one competitor wins a customer

The unit of competition in our market is a **workload**

We compete one workload at a time, and are never fully shut out of an account

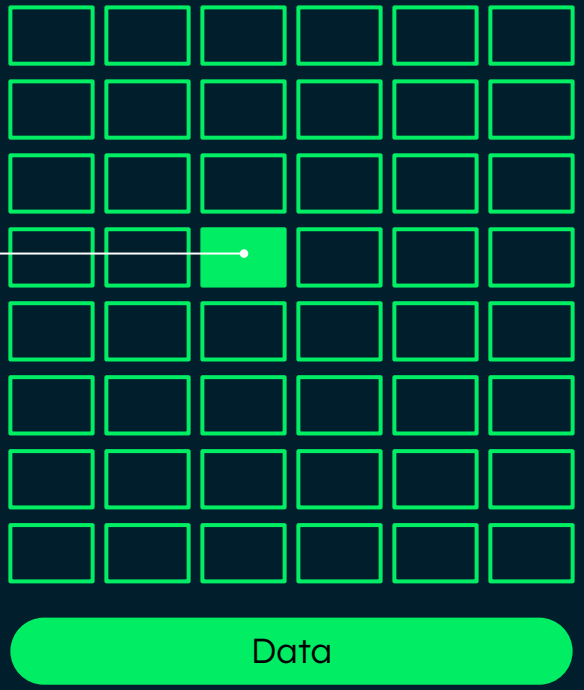

To grow within an account,
we need to acquire more workloads
over time

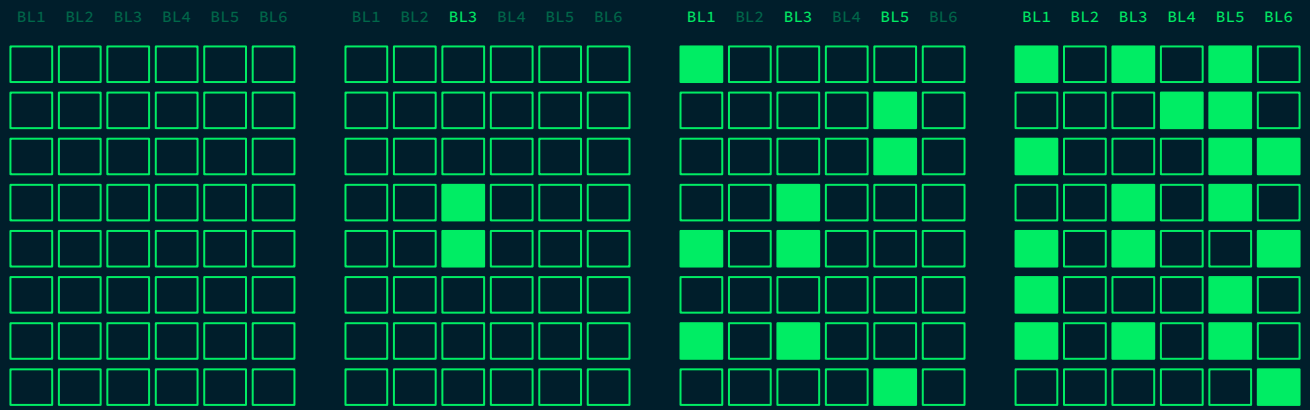
BL1 BL2 BL3 BL4 BL5 BL6
BL1 BL2 BL3 BL4 BL5 BL6
BL1 BL2 BL3 BL4 BL5 BL6
BL1 BL2 BL3 BL4 BL5 BL6

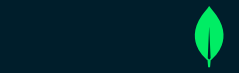

The best way to add more workloads faster is to become a standard in the account

Becoming a standard means that MongoDB is approved to be broadly used in the account

Benefits of becoming a standard

Explicit C-level Support

Company-wide developer access

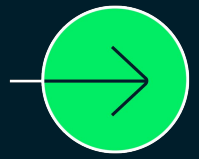

Streamlined app deployment

Benefits of Becoming a Standard

A F100 Financial Institution

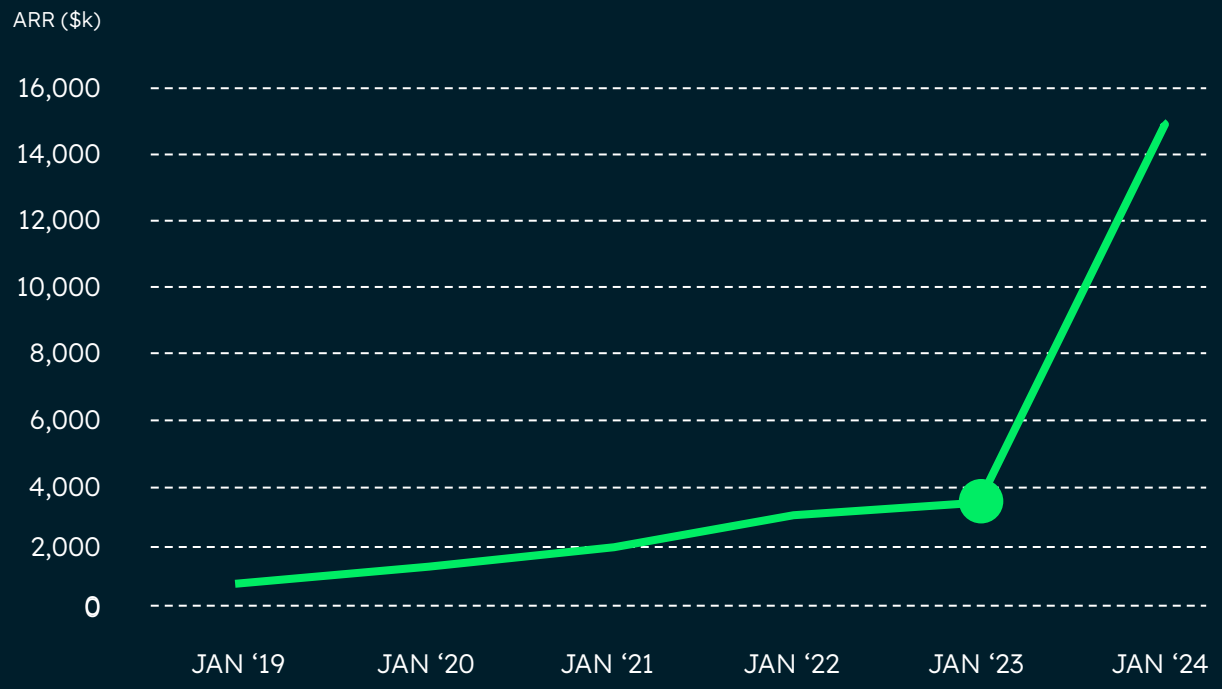

Benefits of Becoming a Standard

A F100 Healthcare Multinational

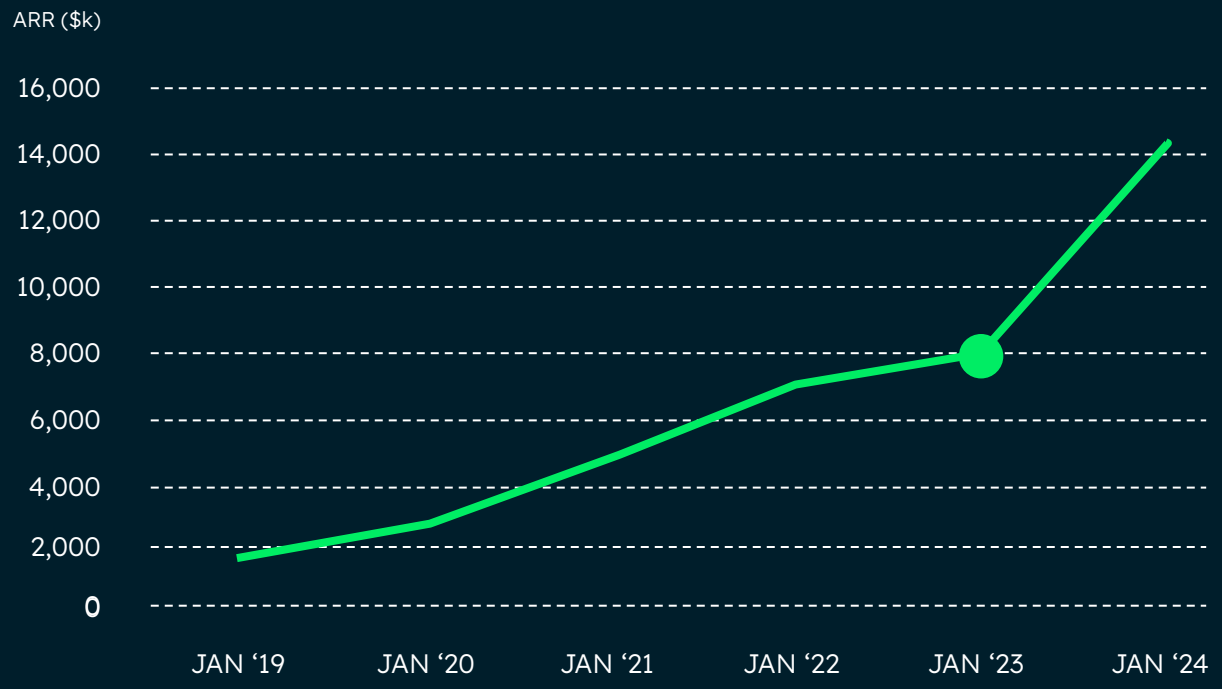

Benefits of Becoming a Standard

A Global Retail Company

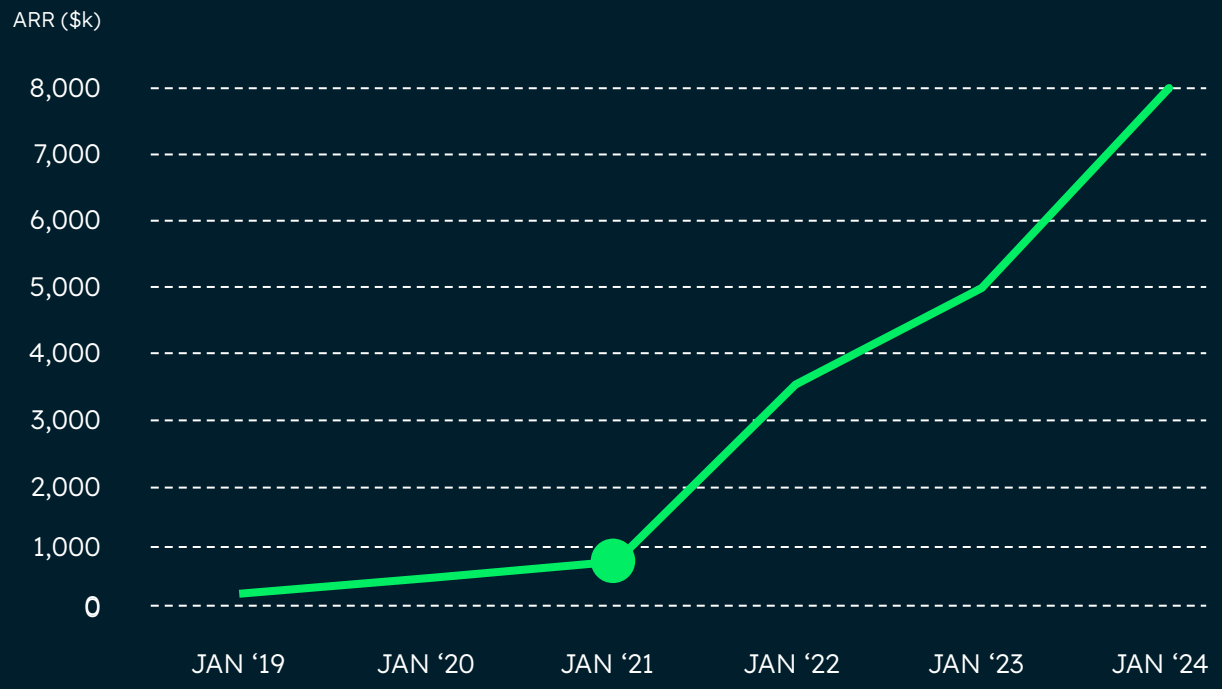

How we become a standard

Solve a broad range of problems

Accelerate application modernization

Establish ourselves as a trusted AI partner

How we become a standard

Solve a broad range of problems

Accelerate application modernization

Establish ourselves as a trusted AI partner

Over the past decade, enterprises have sought to compensate for limitations of legacy technology by introducing niche point solutions into their data estate.

Resulting in architectures that look like this...

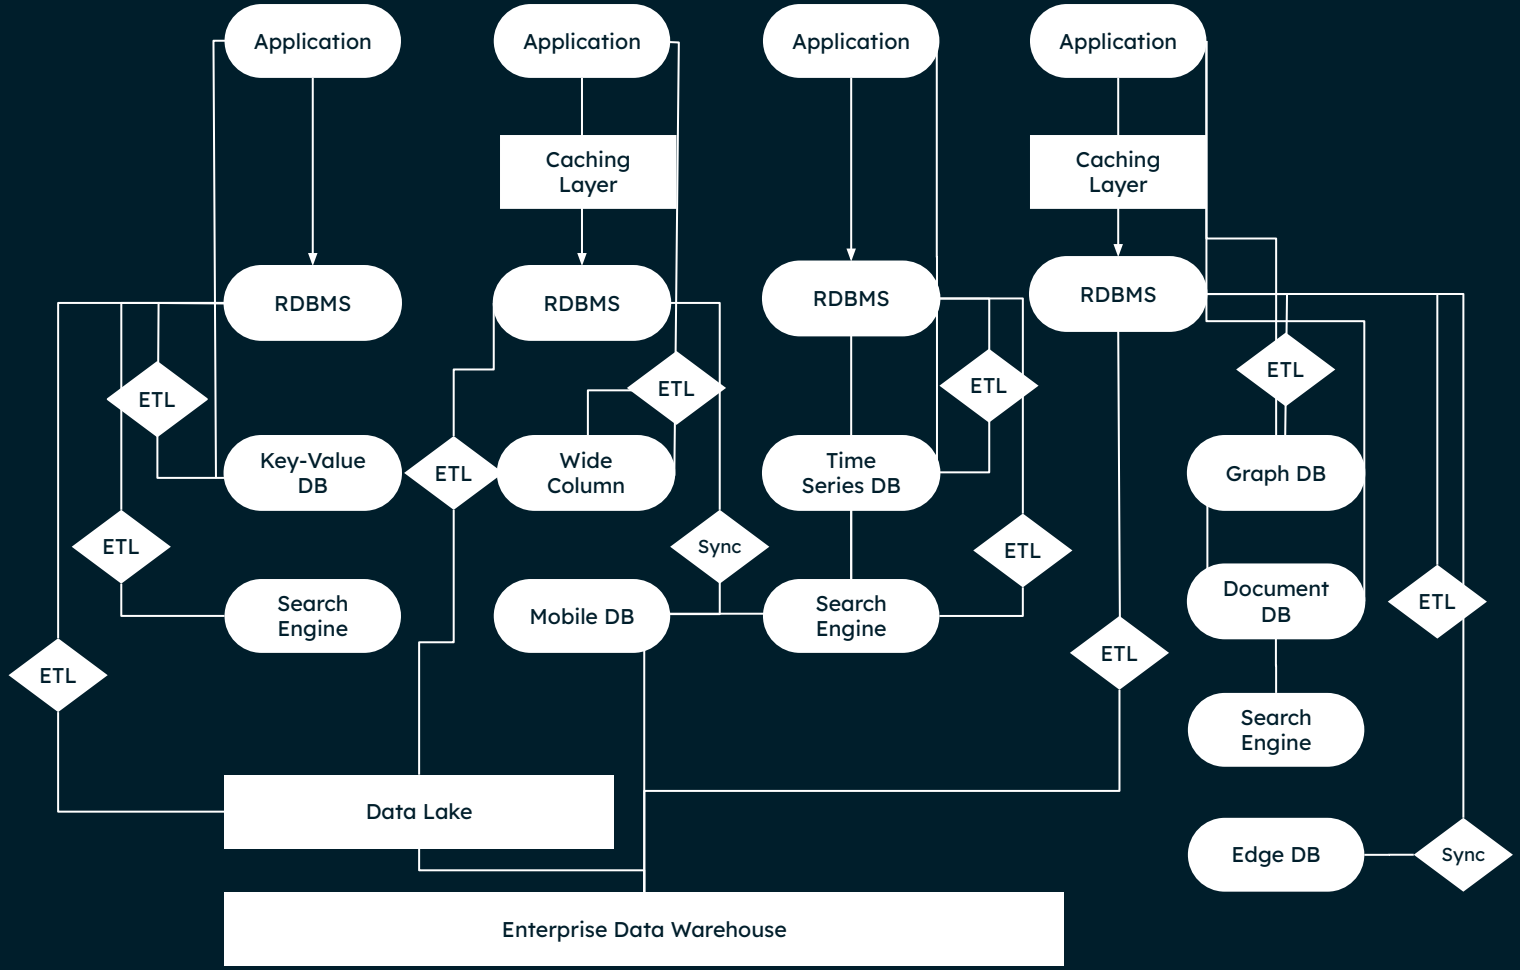
Application
Application
Application
Application
Caching Layer
Caching Layer
RDBMS
RDBMS
RDBMS
RDBMS
ETL
ETL
ETL
ETL
ETL
Key-Value DB
Wide Column
Time Series DB
Graph DB
ETL
Sync
ETL
ETL
Search Engine
Mobile DB
Search Engine
Document DB
ETL
ETL
ETL
Search Engine
Data Lake
Edge DB
Sync
Enterprise Data Warehouse

This complexity inhibits innovation

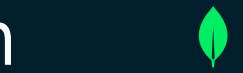

87%

of CIOs agree that complexity is slowing them down

McKinsey: CIO Agenda

75%

of Developers say working with data is the hardest part

MongoDB: Report on Data and Innovation

Enterprises are moving away from point solutions

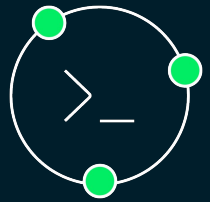

Improve developer productivity

Eliminate back end complexity

Reduce costs and improve ROI

MongoDB's Developer Data Platform is designed to allow customers to address a broad set of use cases across a variety of deployment models on one standard platform

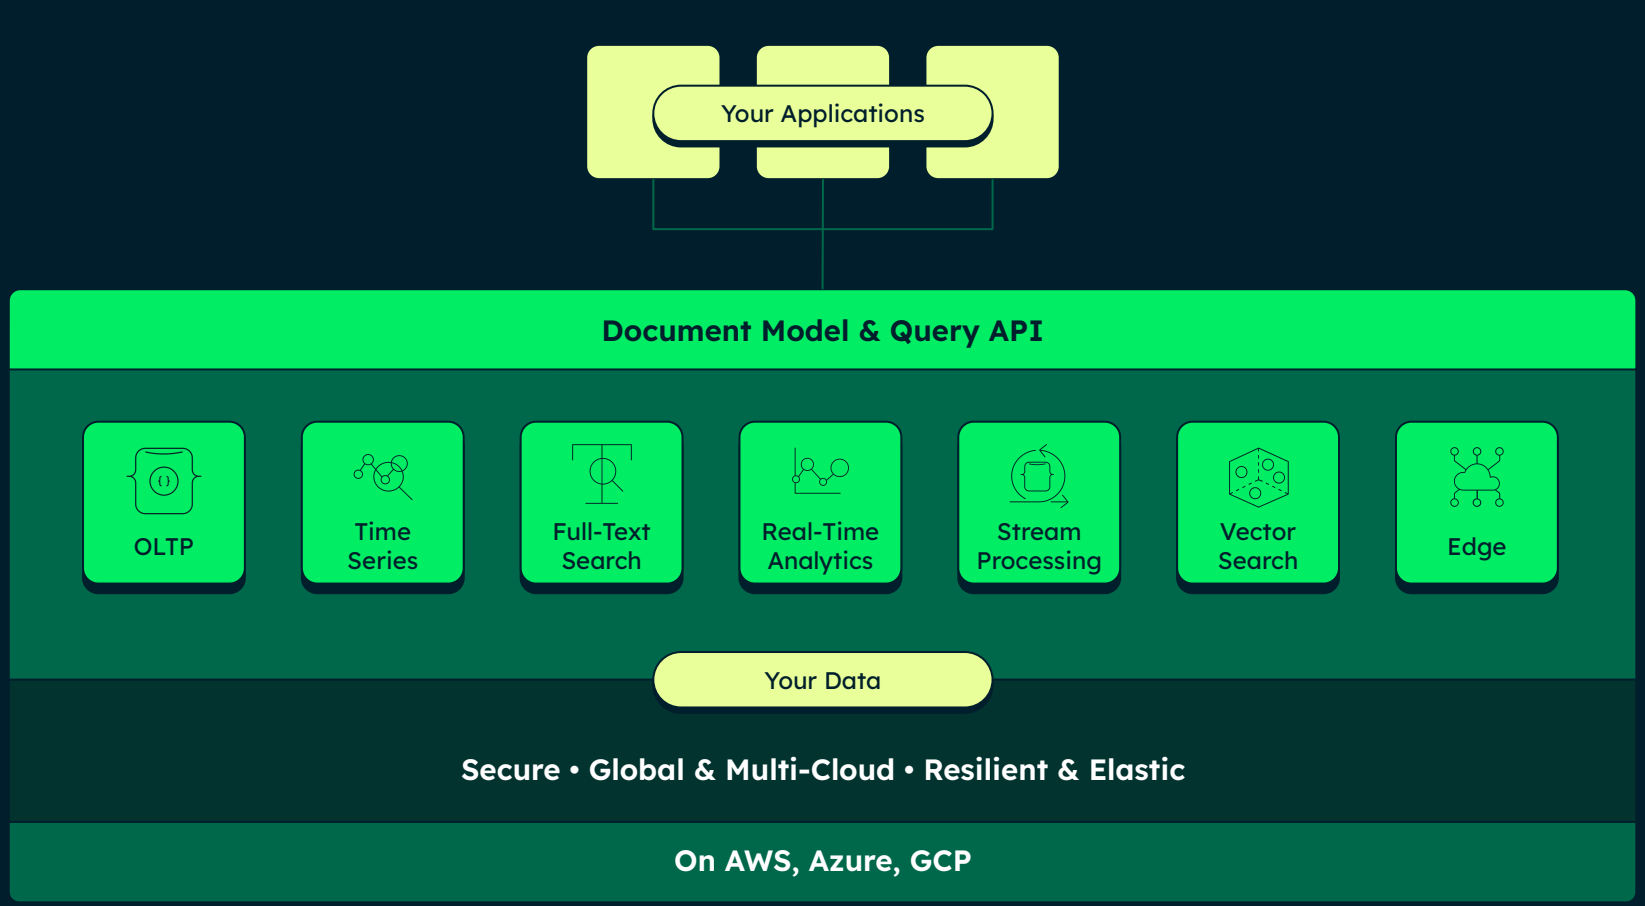
Your Applications
Document Model & Query API
OLTP
Time Series
Full-Text Search
Real-Time Analytics
Stream Processing
Vector Search
Edge
Your Data
Secure • Global & Multi-Cloud • Resilient & Elastic
On AWS, Azure, GCP


{}
OLTP

Reminder: MongoDB's origin story was to address the limitations of relational databases

Rigid to change + imposes unnecessary constraints for developers

Don't cope well with high-fidelity data

Not appropriate for Internet scale

Not appropriate for distributed data

Expensive hardware, punitive licensing, cloud lock-in, intrusive audits

MongoDB OLTP Value Proposition

Developer Experience

Security, Resilience & Scalability

Run Anywhere

MongoDB's foundational technical advantage is the document model

Developer Experience: The document model is flexible and maps to how developers think and code

Flexibility: Documents are a superset of all other data models, allowing us to address the vast majority of operational / transactional use cases

Scalable: Documents put data together in a way that is more performant and efficient and allows almost infinite scalability

Document model brings all data types together

Vector embeddings

Unstructured image data

Generated text descriptions

```
{
   "_id": ObjectId("238478293"),
   "title": "MongoDB TV",
   "description": "All your MongoDB updates, news, videos, and ...",
   "genre": ["Programming", "Database", "MongoDB"],
   “vectorEmbeddings”: [ 0.25, 0.5, 0.75, ... ],
   "seasons": [
      {
         "seasonNumber": 1,
         "episodes": [
            {
               "episodeNumber": 1,
               "thumbnail": “image_65.jpg”,
               "title": "EASY: Build Generative AI Applications",
               "date": ISODate("Oct 5, 2023")
               "description": "Join Jesse Hall….",
            },
            {
              ...
            }
         ]
      },
   ]
}
```

We are seeing success across industries and use cases

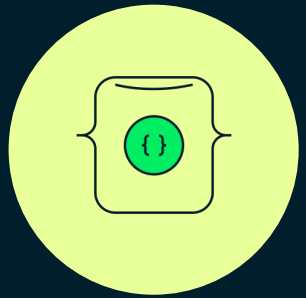

WE ARE FURTHER DRIVING OLTP PERFORMANCE

MongoDB 8.0

- Significant performance improvements (20-60%) across all common workload types, especially common read/write and time series workloads
- Materially enhancing our sharding capabilities, including resharding up to 50x faster and a more balanced distribution of data across shards
- Providing customers better control and performance visibility via Query Insights
- Continue to expand Queryable Encryption functionality, including range support


Full-Text
Search

Search Before Atlas Search

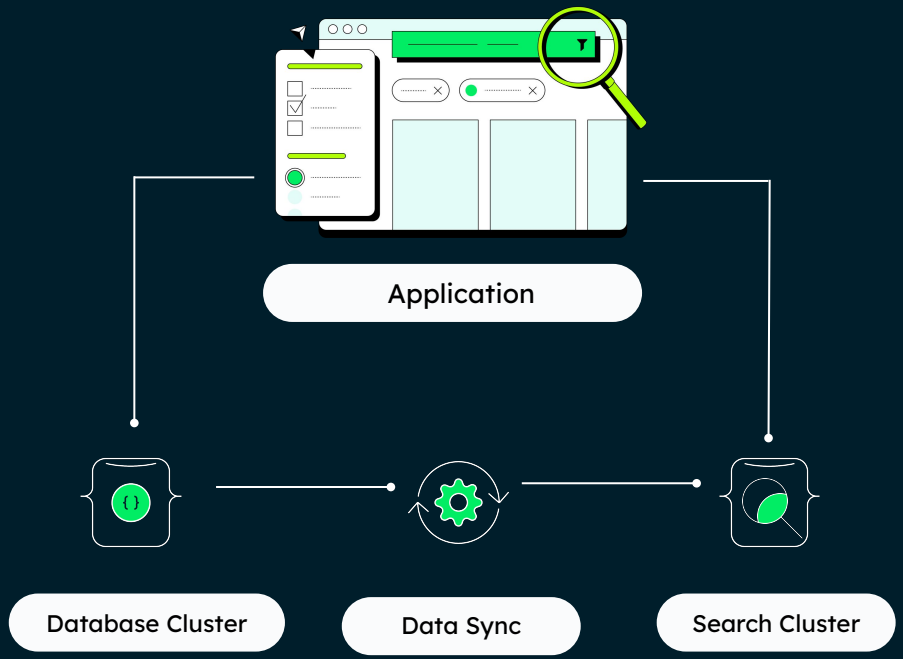

Atlas Search

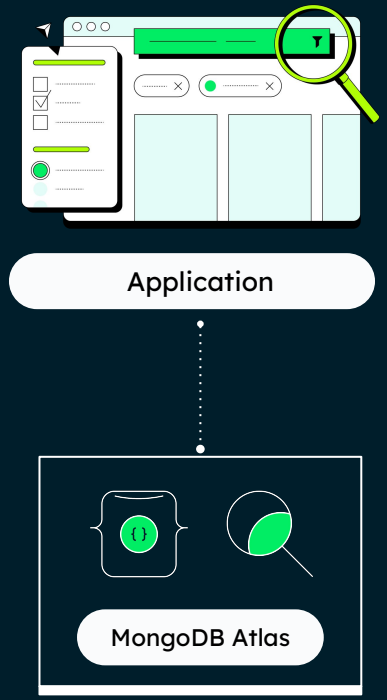

Atlas Search Value Proposition

Eliminate “Synchronization Tax”

Superior Developer Experience

Fully Managed Platform

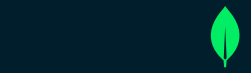

> I thought it was going to be difficult to migrate the data, but it turns out that the toolkit Atlas provides made it incredibly simple. One of the big selling points for Atlas is that it has Atlas Search, meaning I can get rid of our search databases.

— Ken Schuelke,
Division Information Officer
Toyota Financial Services

> MongoDB Atlas Search was a game changer. We ran a proof of concept and discovered how easy it is to use. We can index in one click, and because it's a feature of MongoDB Atlas, we know data is always up-to-date and accurate. We don't need to maintain additional infrastructure as the new features were implemented in less than two weeks, compared to the three months it would have taken to roll out another solution.

— Andrii Hrachov,
Principal Software Engineer
Delivery Hero

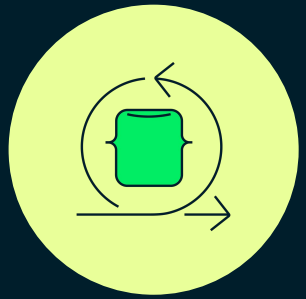

Building a rich Full-Text Search experience

Recent Innovation:

- Dedicated Atlas Search nodes on AWS and GCP
- Support for sequential pagination in Atlas Search
- Atlas Search local experience on Atlas CLI
- Atlas Search query analytics preview

Announcement:

- Dedicated Atlas Search nodes on Azure (Public Preview)
- Search coming to MongoDB Community


Vector
Search

Building AI-enabled apps is difficult

Embedding Generation

The process of representing data as vectors using embeddings models

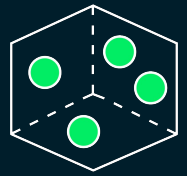

Vector Search

A method for retrieving similar items from a dataset by comparing their vector representations

Database

Contains high value proprietary operational data that can be retrieved to aid an LLM in answering a query

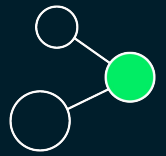

LLMs

Generates output based on relevant data that it receives from the vector search process

Atlas Vector Search Value Proposition

- Streamline RAG with unified data
- Enterprise readiness
- Stay agile with flexible data model

Retool's State of AI 2023 report shows that our Atlas Vector Search value proposition is resonating with customers

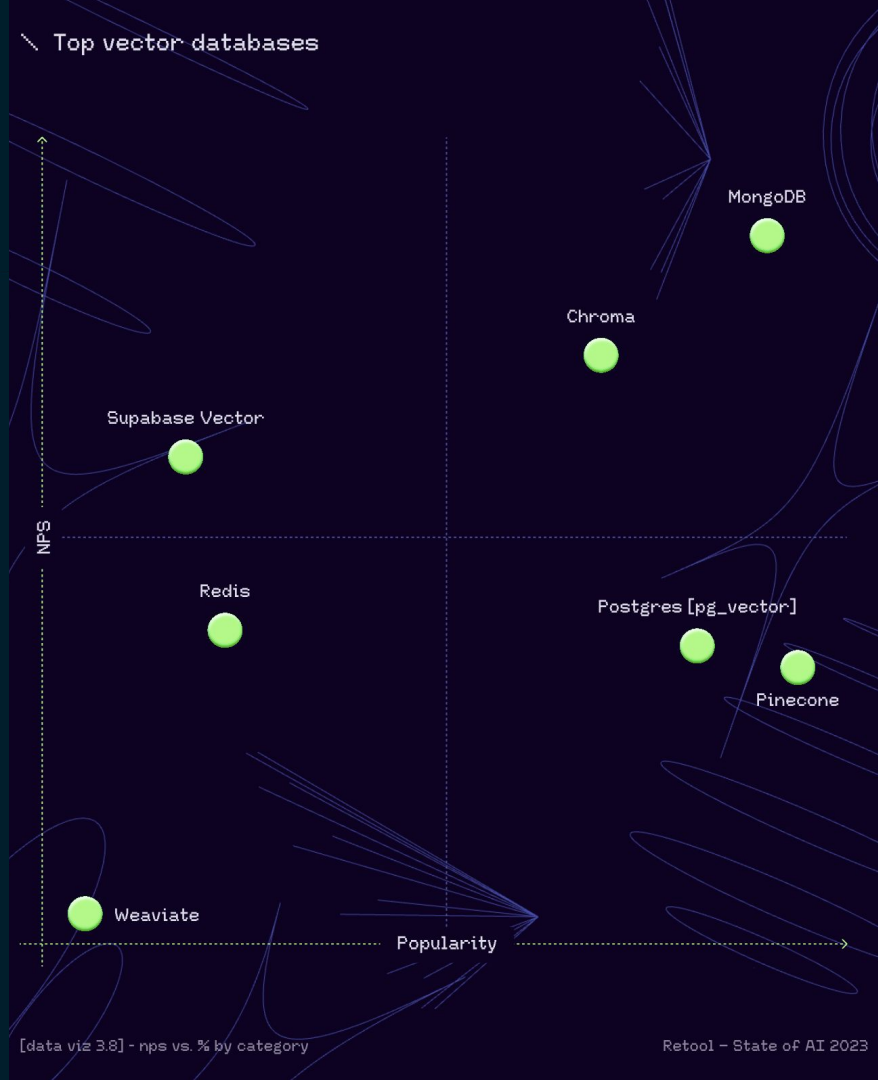

> With Atlas Vector Search, we now possess a battle-tested vector/metadata database, refined over a decade, effectively addressing our dense retrieval requirements. There's no need to deploy a new database, as our vectors and artifact metadata can be seamlessly stored alongside each other.

— Russell Sherman,
Co-Founder & CTO
VISO Trust

> The generative AI we've introduced currently creates vector embeddings from documents, so when a user asks a question, it retrieves the most relevant document and uses LLMs to build the answer. We're looking at migrating vector embeddings into MongoDB Atlas to create a fully integrated, functional system. We'll then be able to use Atlas Vector Search to build AI-powered experiences without leaving the Atlas platform – a much better experience for developers.

— Sabato Severino,
Senior AI Solution Architect
Eni

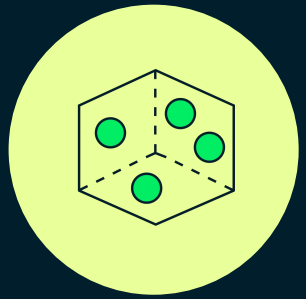

Rapid pace of innovation in Vector Search

Recent Innovation:

- Vector Search local experience with Atlas CLI (Sept 2023)
- Dedicated $vectorSearch Aggregation Stage (Sept 2023)
- Vector Search GA (Dec 2023)
- Create and manage Atlas Vector Search indexes in Compass (Feb 2024)

Announcement:

- Dedicated Vector Search nodes on Azure in Public preview
- Vector Search as a datastore on Haystack and DocArray
- Vector Search as a knowledge base on AWS Bedrock
- Vector Search coming to Community


Stream
Processing

Streaming data is critical for real time use cases

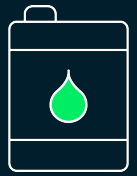

Energy

Energy company rapidly adding IoT devices seeks a **continuous view** into grid health

Manufacturing

Pump Supplier using **continuous monitoring** for pump yield optimization, early failure detection, and equipment health

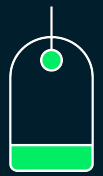

Technology

An engineering software leader wants to offer more **tailored app experiences**

Airlines

A major airline to power a crew maintenance application tracking parts use with **low-latency**

Components of a streaming system

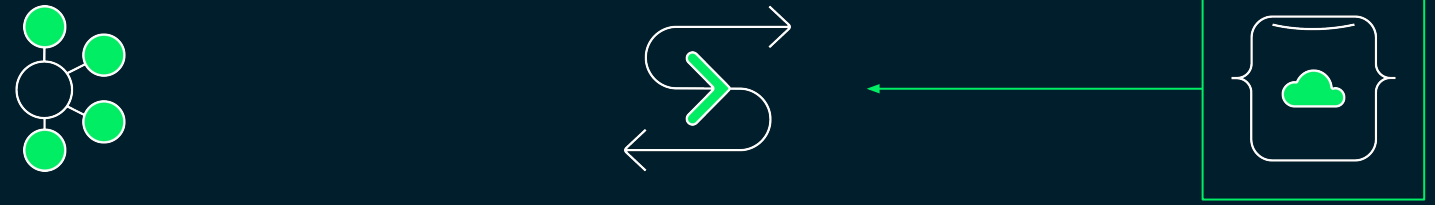

Streaming Transport

Services that ingest data from “sources” as it’s created and then stream that data to targets (“sinks”). Kafka is the dominant technology

Stream Processing

Services that perform real-time queries on streaming data, such as transformations and aggregations

Databases

Ingest and index data from streaming platforms and processors into a persistence layer, making it queryable

Atlas Stream Processing Value Proposition

Data model natural fit for stream processing

An integrated and unified experience

Operational simplicity

> We process about half a billion messages per month in one of our busiest microservices. Atlas Stream Processing has the potential to help us greatly simplify our architecture and leverage powerful operators like windowing and validation are critical to us as our business evolves and scales. It will potentially reduce the latency in our aggregation pipelines.

— Ray Chew,
Principal Software Engineer
Playvox

> Our Acoustic Connect platform must be able to efficiently process and manage millions of marketing, behavioral, and customer signals as they occur. With Atlas Stream Processing, our engineers can leverage the skills they already have from working with data in Atlas to process new data continuously, ensuring our customers have access to real-time customer insights.

— John Riewerts,
EVP Engineering
Acoustic

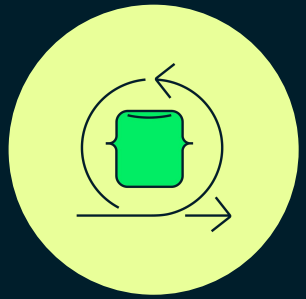

WE ARE EXCITED FOR THE NEXT MILESTONE

Atlas Stream Processing GA

- Initial support for AWS expanding to Azure and GCP
- Integrates with Apache Kafka, Confluent, AWS MSK and other Kafka-compatible streaming platforms such as Azure Event Hubs and Redpanda
- Available presized for Development and Production workloads

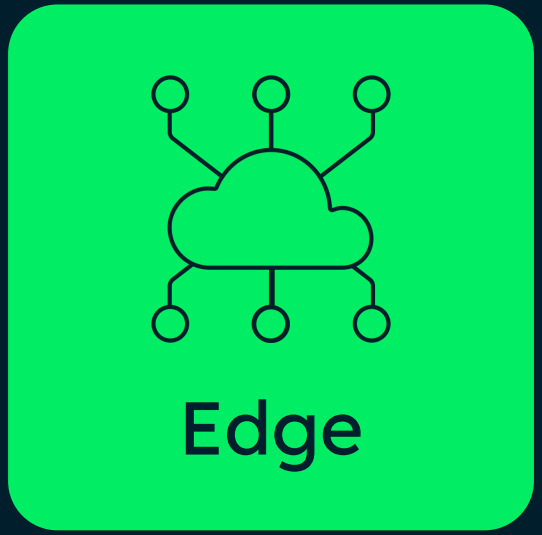
Edge

Cloud computing vs. Edge Computing

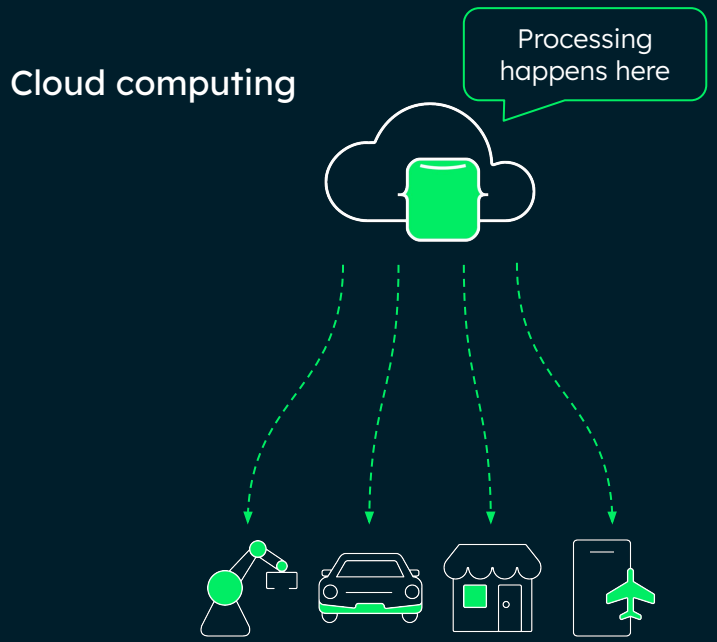

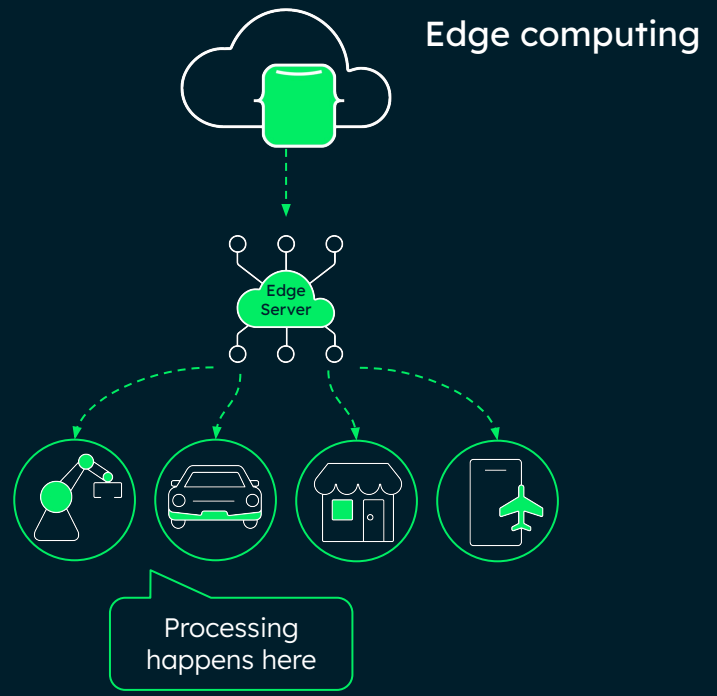

Benefits of Edge Computing

- Local processing and low latency
- Tolerance to disconnection or unreliable connection
- Real-time data enabling time-sensitive use cases
- High volume and variance
- Cost optimization
- Data privacy

Despite the benefits, building edge applications is challenging

Complex data synchronization & management

Diverse hardware and software requirements

Security vulnerabilities and risks

Atlas Edge Server enables end-to-end connectivity

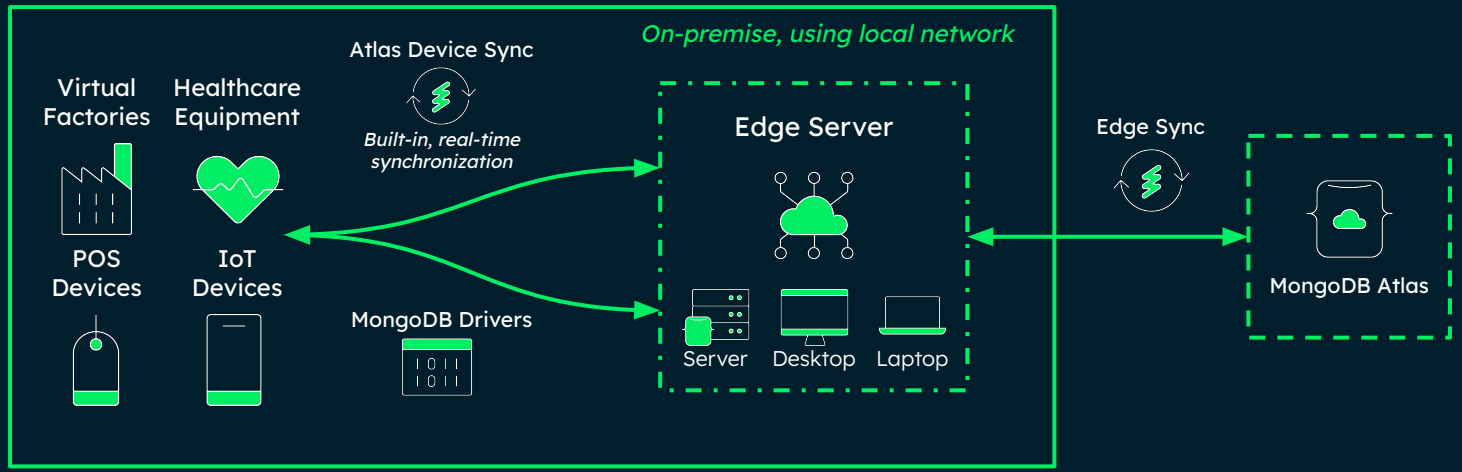

Always-on experience

Collaborative & real-time data

Secure by default

Atlas Edge Server Value Proposition

- Real-time sync
- Optimized bandwidth
- Low latency/always on

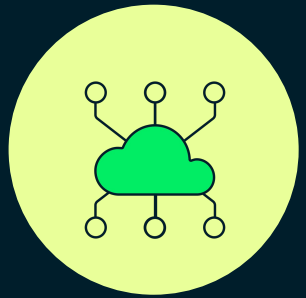

WE ARE OPENING ACCESS TO MORE CUSTOMERS

Atlas Edge Server Public Preview

- Announced Atlas Edge Server Private Preview October 2023
- Local MongoDB paired with a sync server that runs on-premises to enable bidirectional sync between the edge server and Atlas.
- Ensures local capabilities, such as inventory management and point-of-sale, run smoothly in disconnected and offline situations.

How we become a standard

Solve a broad range of problems

Accelerate application modernization

Establish ourselves as a trusted AI partner

Personal Intro: 28 years of data and middleware experience

Paul Done
Field Chief Technology Officer
Modernization Factory

Our principal competitor remains legacy relational technology

Source: IDC, Data Management Software Market, $Bn

Relational technology is no longer capable of meeting the needs of today's modern applications

Rigid to change + imposes unnecessary constraints for developers

Don't cope well with high-fidelity data

Not appropriate for Internet scale

Not appropriate for distributed data

Expensive hardware, punitive licensing, cloud lock-in, intrusive audits

So why does relational technology still have dominant market share?

Because getting off relational is very hard.

There is tremendous variety within relational applications

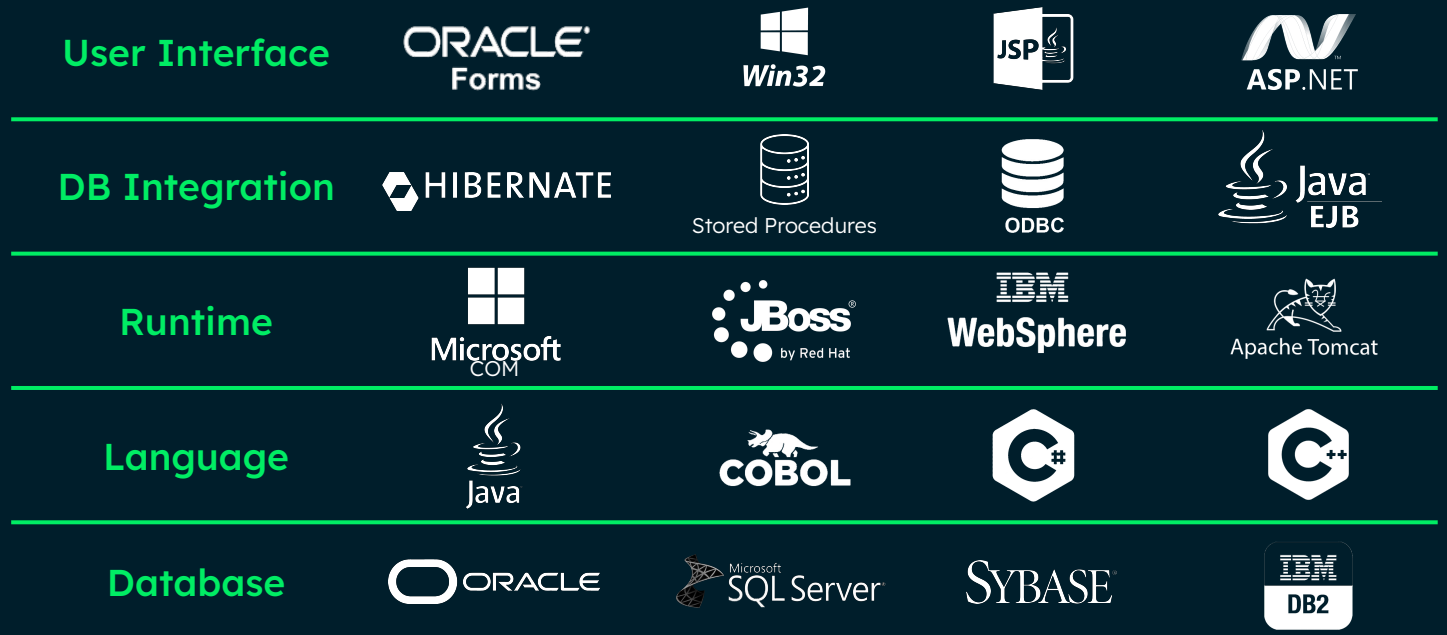

Reminder: three steps in app modernization

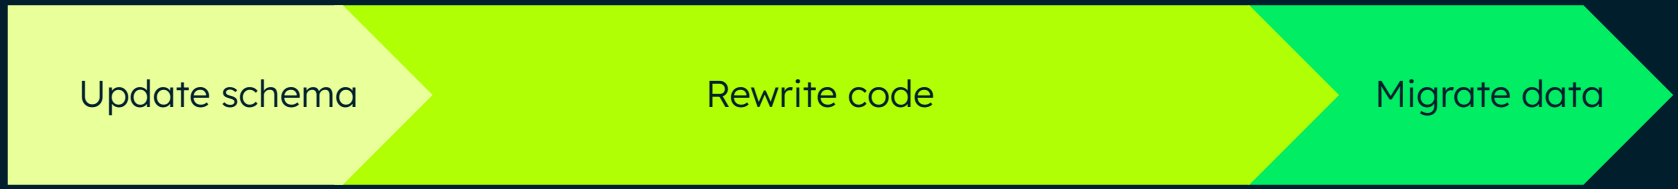

A more detailed look at the process

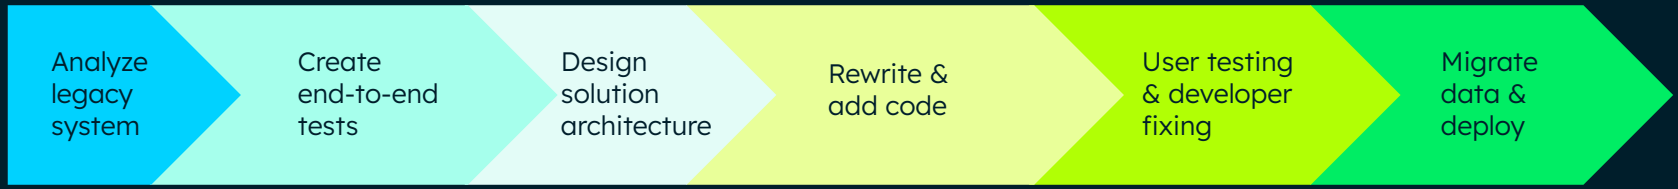

Last year, we announced the general availability of Relational Migrator, a product that simplifies modernization

Relational Migrator makes these steps easier today:

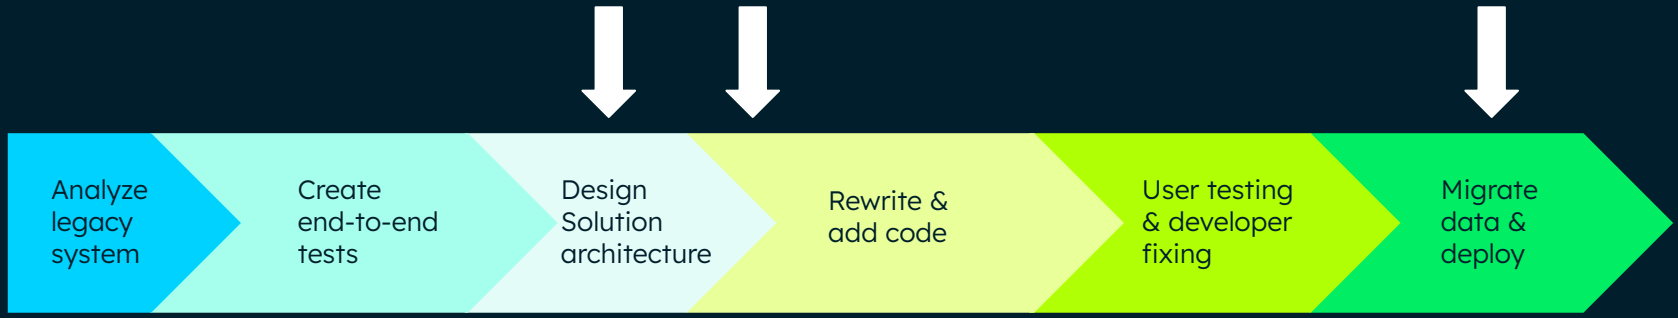

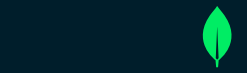

> Relational Migrator is an easy and seamless tool for apps looking to migrate from RDBMS to MongoDB. All we needed to give is the MongoDB destination and with one click the records started migrating to the destination.

— DBA
from a Fortune 50
Health Insurance Company

> Relational Migrator is a great tool to support the migration, especially the graphical support for modeling the data streams from the source structure to the target structure. The MongoDB enabling team did an excellent job in helping us to be able to go through the migration steps.

— Lead Architect
from a Global 50
Financial Services Company

But we can make app modernization much easier thanks to AI

We are running several pilots this year to learn how much easier

AI can simplify every step of app modernization

Activity	AI Acceleration	AI Impact
ANALYZE LEGACY SYSTEM	Use LLMs to analyze legacy codebase in order to understand its structure, eliminating the need for input from original developers who will have moved on	✔✔
CREATE END-TO-END TESTS	Analyze how users interact with the app and then feed that information to an LLM to create tests that ensure the viability of the new application	✔✔✔
DESIGN SOLUTION ARCHITECTURE	Identify how to "pull out" potential microservices from the monolith by using existing source code + recorded end-user behaviour with LLMs	✔
REWRITE & ADD CODE	Partially automate the creation of microservices by using existing source code + recorded end-user behaviour with LLMs	✔✔✔
USER-TESTING & DEVELOPER-FIXING	Human intervention for testing and bug code fixing is still needed but can be streamlined through the use of LLMs	✔✔
MIGRATE DATA & DEPLOY APP	Reduce deployment risk by using the LLMs to identify how best to structure code to dual-run old and new databases during intermediate phases of migration	✔

Some (very early) lessons from the pilots

AI-Generated tests hold great promise to shorten development and testing cycles and reduce risk

Capabilities of LLMs vary significantly, increasing the importance of picking right

“Last mile” human expertise will remain critical to validate code and ensure quality

Where do we go from here?

FY 24

FUTURE

Pilots

This year we are focused on deep engagements with a small number of customers, to prototype and iterate on a number of approaches for AI-enhanced app modernizations

Playbooks/Templates

In the future, we'll consolidate the learnings to build processes and templates to accelerate app modernization for each customer and also scale the number of engagements we can run

Product

We expect that app modernization will always involve some human interaction. However, over time, we will incorporate many of our learnings into Relational Migrator, the product.

In summary...

Getting off relational is very hard

We are learning how to harness the power of AI to make the process much easier

How we become a standard

Solve a broad range of problems

Accelerate application modernization

Establish ourselves as a trusted AI partner

Enterprise customers tell us that AI strategy is a board-level priority, due to its promise to dramatically change cost structures and enhance customer experiences

Customers also tell us that they are struggling to formulate and execute a strategy because the AI space is complex and fast-moving

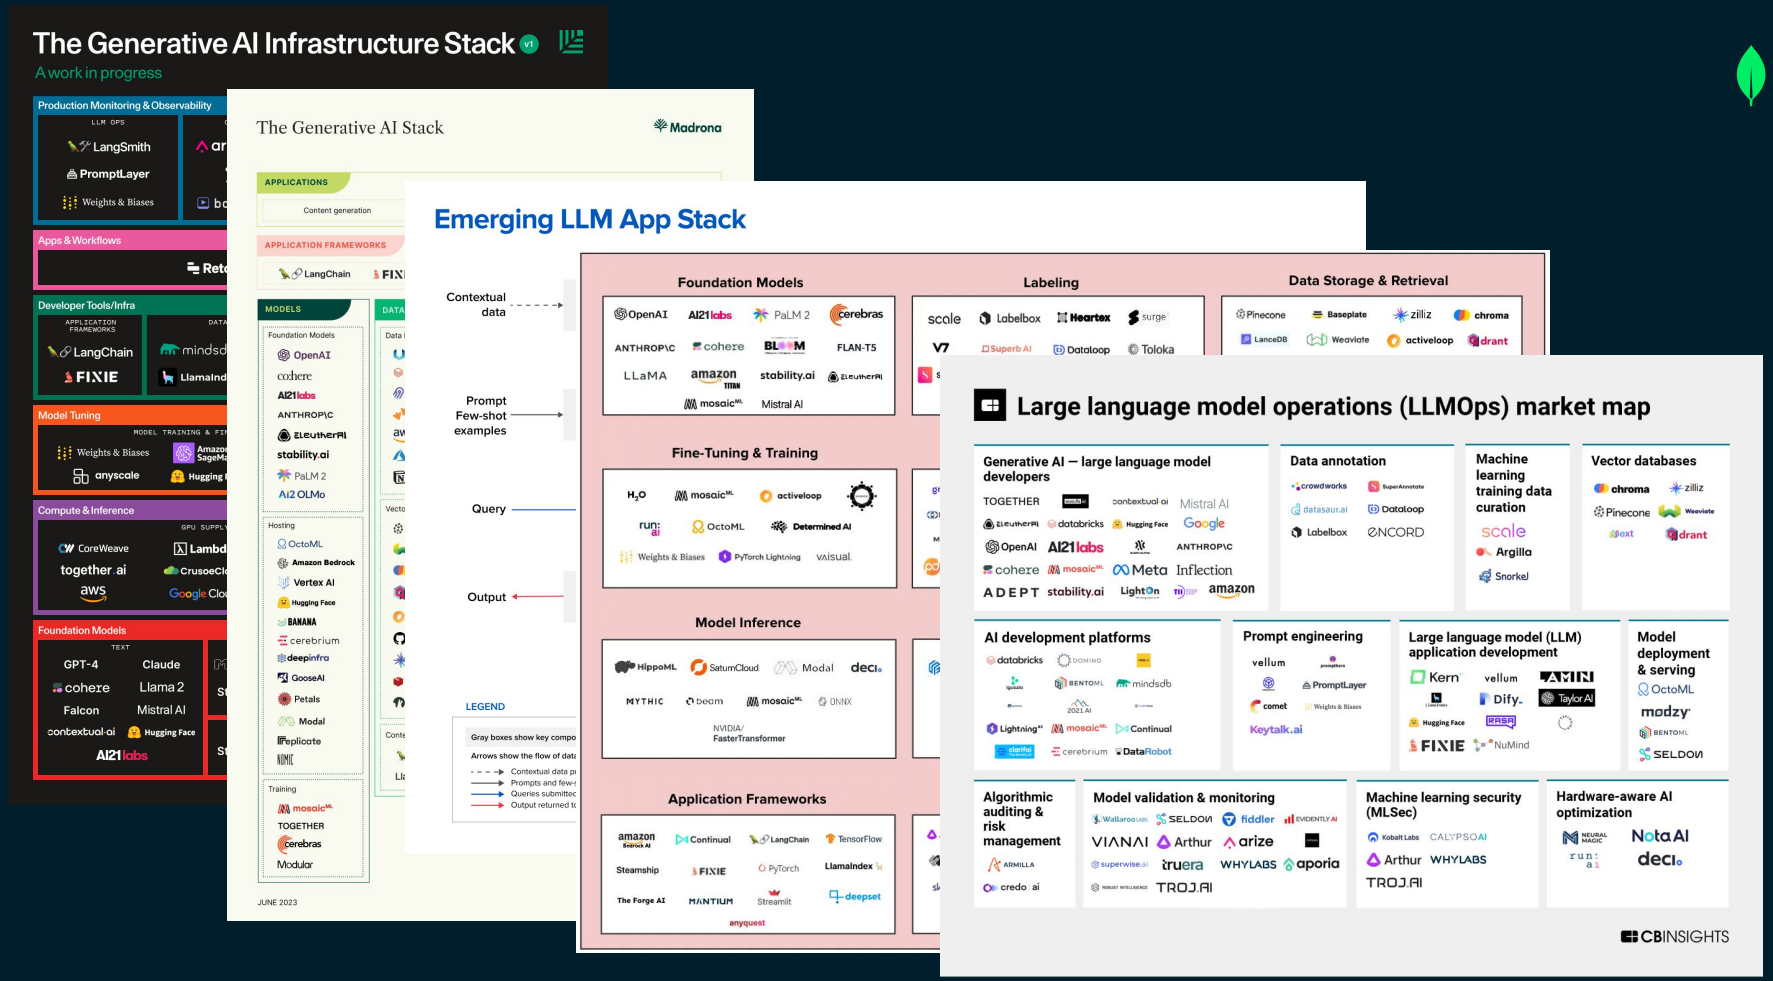
The Generative AI Infrastructure Stack
v1
A work in progress
Production Monitoring & Observability
LLM OPS
LangSmith
PromptLayer
Weights & Biases
Apps & Workflows
Developer Tools/Infra
APPLICATION FRAMEWORKS
LangChain
FIXIE
Model Tuning
MODEL TRAINING & FI
Weights & Biases
anyscale
Compute & Inference
GPU SUPPL
CoreWeave
together.ai
aws
Foundation Models
TEXT
GPT-4
Claude
cohere
Llama 2
Falcon
Mistral AI
contextual.ai
Hugging Face
AI21labs
The Generative AI Stack
Madrona
APPLICATIONS
Content generation
APPLICATION FRAMEWORKS
LangChain
MODELS
Foundation Models
OpenAI
cohere
AI21labs
ANTHROPIC
stability.ai
PaLM 2
AI2 OLMo
Hosting
OctoML
Amazon Bedrock
Vertex AI
Hugging Face
BANANA
cerebrium
deepinfra
GooseAI
Petals
Modal
Replicate
Training
TOGETHER
cerebras
Modular
JUNE 2023
Emerging LLM App Stack
Contextual data
Prompt Few-shot examples
Query
Output
LEGEND
Foundation Models
OpenAI
AI21labs
PaLM 2
cerebras
ANTHROPIC
cohere
BLOOM
FLAN-T5
LLaMA
amazon TITAN
stability.ai
Mistral AI
Labeling
scale
Labelbox
Heartex
surge
Superb AI
Dataloop
Toloka
Data Storage & Retrieval
Pinecone
Baseplate
zilliz
chroma
LanceDB
Weaviate
activeloop
drant
Fine-Tuning & Training
H2O
activeloop
OctoML
Determined AI
Weights & Biases
PyTorch Lightning
Model Inference
HippoML
SaturnCloud
Modal
MYTHIC
beam
ONNX
NVIDIA/ FasterTransformer
Application Frameworks
amazon Bedrock AI
Continual
LangChain
TensorFlow
Steamship
FIXIE
PyTorch
LlamaIndex
The Forge AI
MANTIUM
Streamlit
deepset
anyquest
Large language model operations (LLMOps) market map
Generative AI — large language model developers
TOGETHER
contextual.ai
Mistral AI
databricks
Hugging Face
Google
OpenAI
AI21labs
ANTHROPIC
cohere
Meta
Inflection
ADEPT
stability.ai
amazon
Data annotation
crowdworks
SuperAnnotate
datasaur.ai
Dataloop
Labelbox
NCORD
Machine learning training data curation
scale
Argilla
Snorkel
Vector databases
chroma
zilliz
Pinecone
Weaviate
drant
AI development platforms
databricks
DOMINO
BENTOML
mindsdb
2021.AI
Lightning AI
Continual
cerebrium
DataRobot
Prompt engineering
vellum
PromptLayer
comet
Weights & Biases
Keytalk.ai
Large language model (LLM) application development
Kern
vellum
Dify.
Taylor AI
Hugging Face
FIXIE
NuMind
Model deployment & serving
OctoML
modzy
BENTOML
SELDON
Algorithmic auditing & risk management
ARMILLA
credo ai
Model validation & monitoring
Wallaroo LABS
SELDON
fiddler
EVIDENTLY AI
VIANAI
Arthur
arize
superwise.ai
truera
WHYLABS
aporia
TROJ.AI
Machine learning security (MLSec)
Kobalt Labs
CALYPSOAI
Arthur
WHYLABS
TROJ.AI
Hardware-aware AI optimization
NEURAL MAGIC
Nota AI
run:ai
deci
CBINSIGHTS

Introducing the next generation of Claude

Mar 4, 2024 • 7 min read

Introducing Meta Llama 3: The most capable openly available LLM to date

UPDATED 20:11 EDT / APRIL 10 2024

Mistral AI debuts Mixtral 8x22B, one of the most powerful open-source AI models yet

Apr 04, 2024

Introducing Command R+: A Scalable LLM Built for Business

AIDAN GOMEZ

NEWSROOM PRODUCT

15 Apr

Reka Core: Our Frontier Class Multimodal Language Model

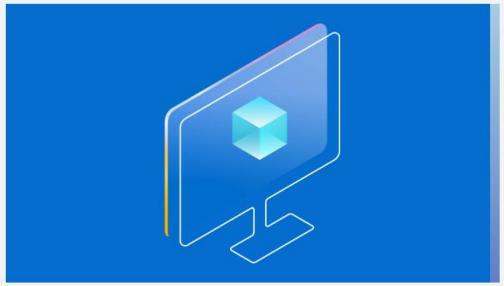

AI + Machine Learning, Announcements, Azure AI, Azure AI Studio

Introducing Phi-3: Redefining what's possible with SLMs

By Misha Bilenko, Corporate Vice President, Microsoft GenAI

OpenAI @OpenAI · Apr 9

Majorly improved GPT-4 Turbo model available now in the API and rolling out in ChatGPT.

Introducing DBRX: A New State-of-the-Art Open LLM

by The Mosaic Research Team

March 27, 2024 in Mosaic AI Research

Share this post

AI/ML 5 Min

Introducing Marengo-2.6: A New State-of-the-Art Video Foundation Model for Any-to-Any Search

This can clearly feel overwhelming

87%

of customers don’t feel equipped
to transform with AI

— WEBER SHANDWICK – UN/PREDICTIONS 2024

Customers do understand they need a holistic and flexible approach to AI

- Any Use Case
- Any Model
- Any Cloud
- Any Data Type

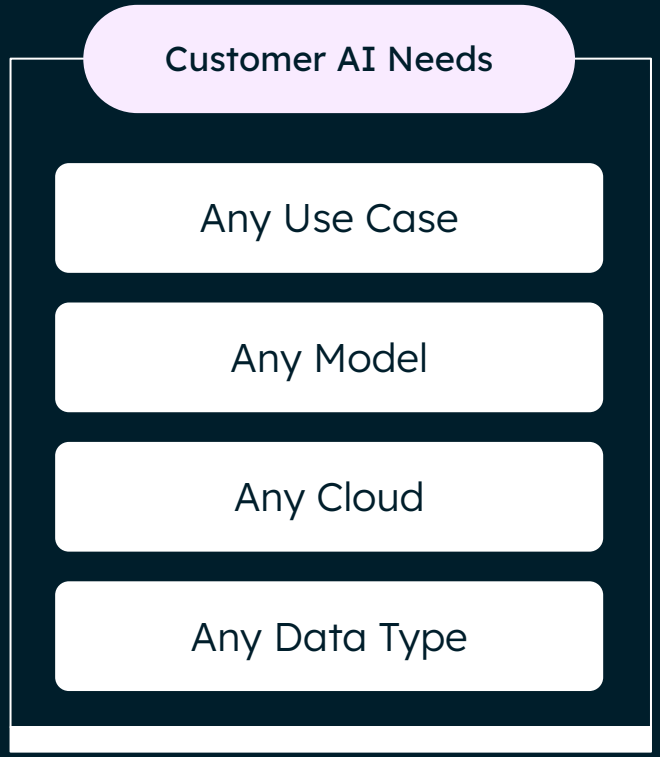
Customer AI Needs
Any Use Case
Any Model
Any Cloud
Any Data Type

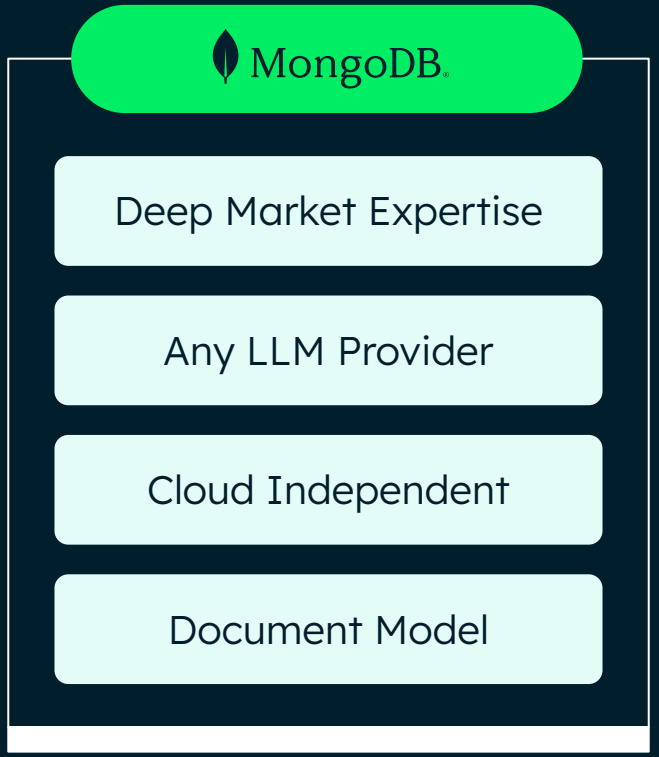
MongoDB.
Deep Market Expertise
Any LLM Provider
Cloud Independent
Document Model

In other words, we are well positioned to serve as a trusted AI advisor to our customers as they navigate this mega trend

But what does that have to do with becoming a standard?

Traditionally, we grow in an account through a land-and-expand motion

We land in an account by engaging with developers to help them solve problems they cannot solve with their existing tech stack

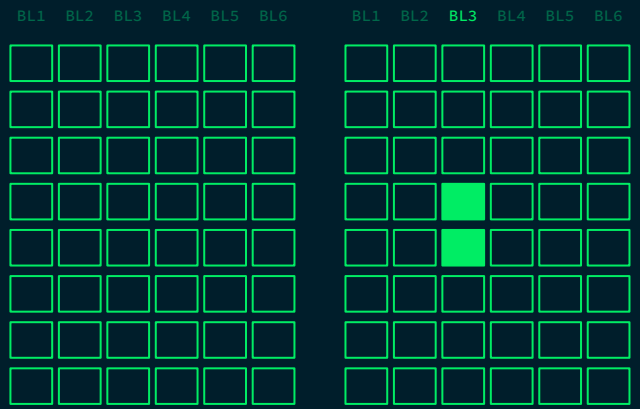

We expand by building momentum with more developer teams and business unit leads

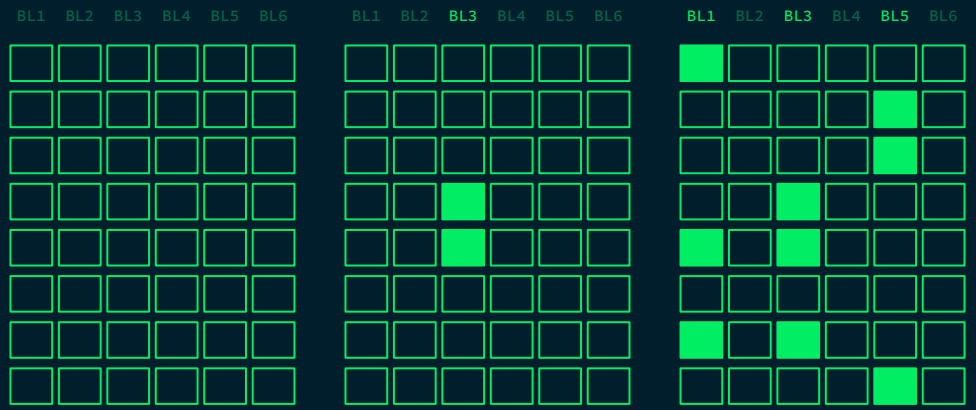

We become a standard when we are so spread in the account that the C-suite realizes the strategic benefits of partnering with us

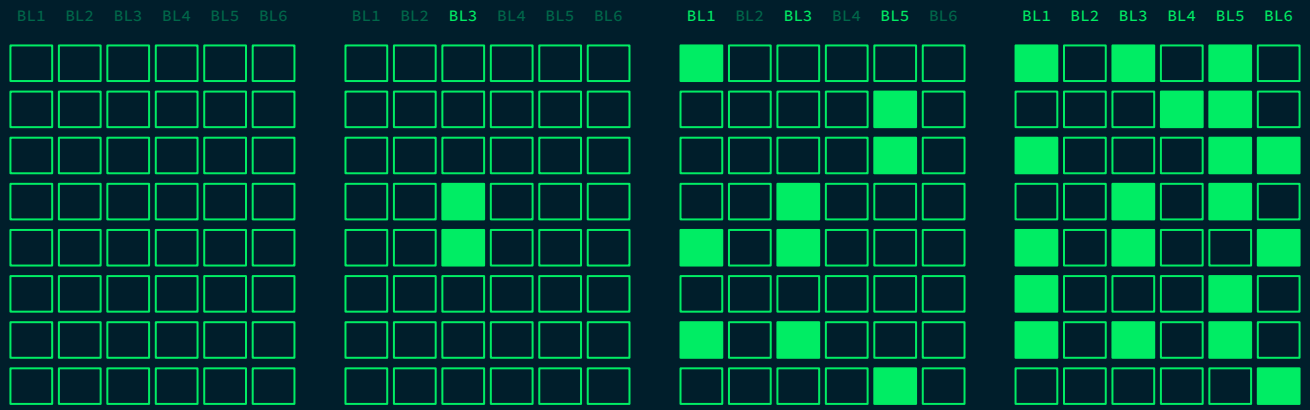

Increasingly, thanks to our AI positioning, we are engaging in strategic C-level conversations with customers sooner, giving us an opportunity to sell tops down to accelerate our land-and-expand motion

Case study: GenAI Executive Briefing

We organized a GenAI executive briefing in London on April 18th

The event attracted 40 C-level execs from 10 countries, representing verticals including finance, retail, technology and government among others

“The depth of discussion was invaluable, MongoDB are uniquely positioned to garner such a cross section of CxOs.”

— Head of AI, Large Global Consulting Firm

How will we seize this AI moment?

We will continue with AI innovation in our product

Infrastructure and Vectors:

- $vectorSearch Aggregation Stage
- MongoDB BSON Specification for Vectors
- Dedicated nodes for Vector Search
- Programmatic deployment for Search nodes (Terraform + Cloudformation)
- Broad Vector Search performance improvements
- Local Development with Atlas CLI
- Vector Search coming to MongoDB Community

AI Stack Integrations:

- New partnerships with Fireworks.ai and Mistral, adding to a growing list
- Atlas Vector Search as an AWS Bedrock Knowledge Base GA
- Dedicated Langchain-MongoDB package in Python and JS
- OpenAI ChatGPT retrieval plugin for MongoDB
- Microsoft Semantic Kernel integration for Atlas Vector Search
- Atlas Vector Search as a datastore on Haystack and DocArray

Productivity & Migrations:

- Collaboration with Amazon CodeWhisperer, Microsoft Copilot and Google Gemini
- Intelligent query experience in Compass
- Vector Search index management in Compass
- Create visualizations using natural language in Charts
- SQL-to-Query API Conversion in Relational Migrator
- Upcoming innovations in application modernization (code analysis, schema mapping, code conversion)

We will keep expanding our AI ecosystem

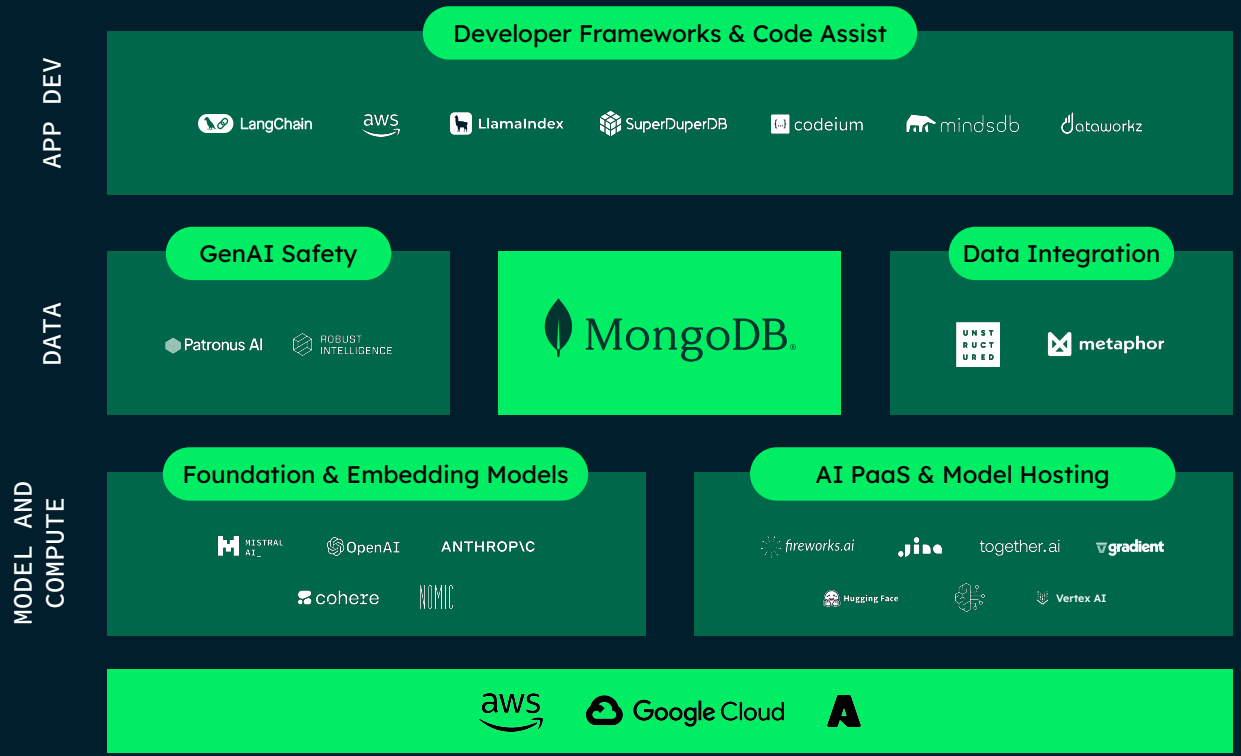

We will make it easier for customers to get started

ATLAS

Pre-configured and optimized MongoDB products, features, and services

PARTNERS

Pre-selected, leading AI tech partnerships

REFERENCE ARCHITECTURES

Proven, repeatable architectures for simple implementation

SERVICES

PS & expert Boutique SIs to handle implementations