
CASE
STUDY

**EPAM HELPS EDMUNDS
MOVE THEIR DATA CLOSER
TO THE BUSINESS WITH A
NoSQL SOLUTION**



CASE STUDY:

EPAM HELPS EDMUNDS MOVE THEIR DATA CLOSER TO THE BUSINESS WITH A NoSQL SOLUTION

Edmunds.com, a long-time EPAM client, arms consumers with the information necessary to “discover, price, and buy the car that is right” for their needs. Providing simply the best, most up-to-date insider data about nearby dealerships, market prices, sales, and car specifications, Edmunds delivers paid advertising for car manufacturers as well as leads- and subscription-based services to dealers. With Edmunds, potential customers get all of the essential information about the car they’re shopping for before they ever step through the door of a dealership.

THE CHALLENGE: FIND AND INTEGRATE NEW OPEN-SOURCE STORAGE SOLUTION

Edmunds approached EPAM to find a solution for replacing the expensive, hard-to-maintain Oracle Coherence in-memory data grid that had been powering their website for five years and simultaneously limiting the company’s platform evolution. The legacy architecture kept the data in one large cluster featuring multiple datasets consisting of 4 to 24 Coherence nodes each, and the whole cluster had 80+ nodes in total. Needless to say, there were a lot of data structures to migrate, which proved to be a challenge in itself.

Edmunds required an open-source, fast-performing NoSQL solution that would organize, embed, and monitor its data based on use cases. Moreover, the client needed a partner that could seamlessly integrate the solution into the existing platform, keeping APIs intact without jeopardizing the performance. EPAM came through as both a solution provider and a partner every step of the way.



CASE STUDY:

EPAM HELPS EDMUNDS MOVE THEIR DATA CLOSER TO THE BUSINESS WITH A NoSQL SOLUTION

THE SOLUTION: A MORE DYNAMIC, LESS EXPENSIVE DATABASE MANAGEMENT SYSTEM

The first step in finding the proper solution for Edmunds was the launch of a joint investigation into which database would best fit the client's needs. This resulted in the NoSQL market analysis phase, during which EPAM analyzed roughly 15 different data storages characterized as distributed in-memory data grids, document-oriented databases, and key value databases. EPAM then ranked the storages by criteria such as licensing and operational costs, performance, applicability, reliability, ease of maintenance, and others.

After conducting comprehensive research on each storage solution, EPAM chose three candidates for use in proof-of-concept testing. Our team started to run multiple learning loops in order to choose the best candidate to replace Coherence and, based on the results, chose MongoDB as the most balanced solution for the client.

In order to meet strict performance requirements throughout the migration, we applied two types of data architecture as follows:

PREFERRED SOLUTION:

MICROSERVICE ARCHITECTURE

Pros:

- Preferable for almost all datasets except the biggest ones
- Capable of clearly specifying all data usages and response formats
- Single point of control

AUXILIARY SOLUTION:

DIRECT DATA ACCESS ARCHITECTURE

Pros:

- Used for datasets with huge amounts of use cases and/or data to transfer
- Better overall performance
- Allows use of additional features such as lazy loading

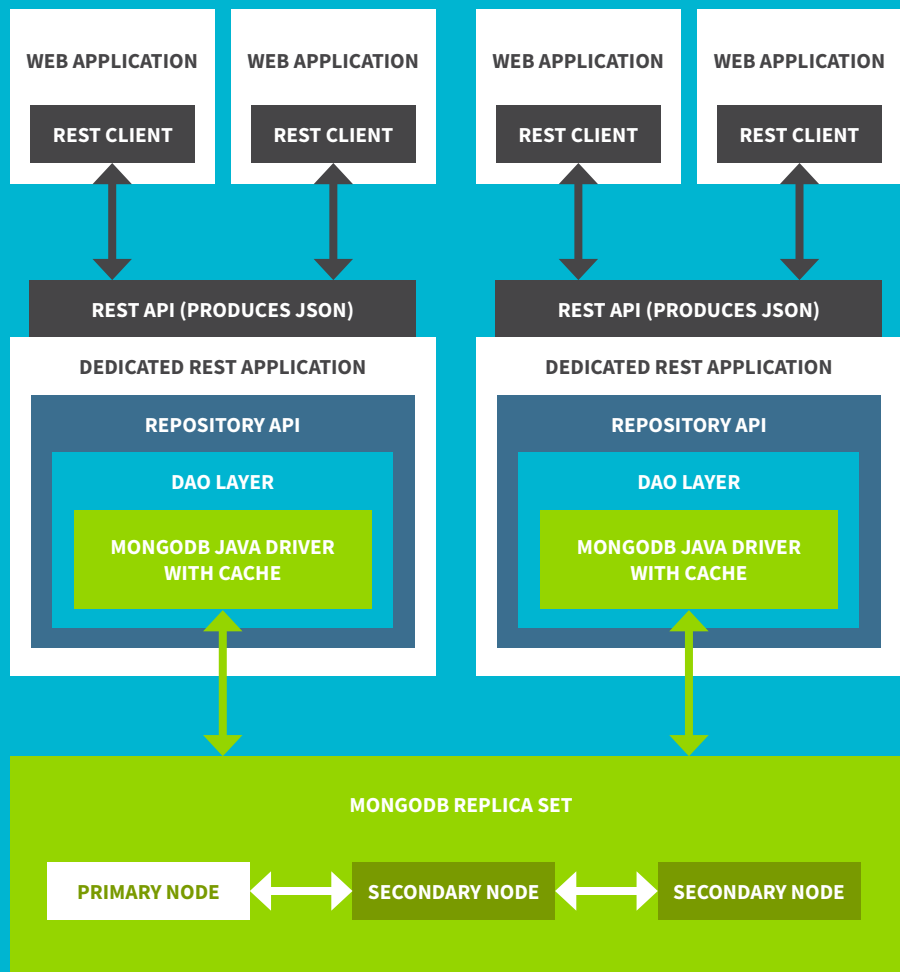
Cons:

- More vulnerable to incorrect data usage
- No single point of control

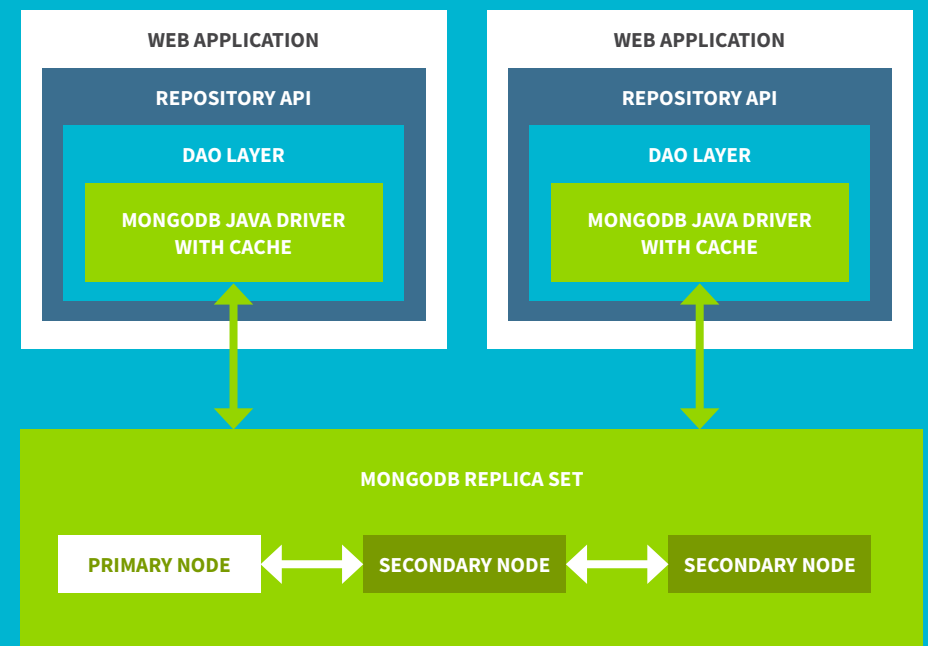
View an in-depth comparison of the two architectures on the next page.

MICROSERVICE ARCHITECTURE VS. DIRECT DATA ACCESS ARCHITECTURE

MICROSERVICE ARCHITECTURE



DIRECT DATA ACCESS ARCHITECTURE





CASE STUDY:

EPAM HELPS EDMUNDS MOVE THEIR DATA CLOSER TO THE BUSINESS WITH A NoSQL SOLUTION

With the right storage solution selected and the data architectures in place, the EPAM team was ready to begin the migration of Edmunds' core datasets. We seamlessly migrated the following datasets to MongoDB:

- **Vehicle:** a configurations, specifications, and pricing dataset
- **Region:** a DMA, geo regions, and incentive region dataset
- **Incentive:** a manufacturer incentive and rebate dataset
- **TMV:** a dataset that powers the True Market Value calculator
- **TCO:** a dataset that powers the True Cost to Own calculator for new and used cars
- **Partner:** information about Edmunds partners and contracts
- **Rules:** rules for Dealer and Inventory data processing
- **Maintenance:** car maintenance actions, recalls, and service bulletins dataset

Throughout the migration process, our team worked diligently to ensure that Edmunds' front-end applications and internal tools would be able to access the data via the shared REST and Java APIs. We also implemented a comparison tool to measure performance of numerous UI pages and REST responses across different environments to make sure that performance was not affected while moving from Coherence to MongoDB.



CASE STUDY:

EPAM HELPS EDMUNDS MOVE THEIR DATA CLOSER TO THE BUSINESS WITH A NoSQL SOLUTION

THE RESULT: AN EXTREME PERFORMANCE UPGRADE FOR EDMUNDS

After one-and-a-half years of hard work and collaboration with Edmunds on the dataset migration project, EPAM successfully completed the replacement of Oracle Coherence with MongoDB in December 2015. This project is another in EPAM's very impressive series of collaborations in partnership with Edmunds – a partnership made stronger by EPAM's ability to accomplish the following:

- Refactor the client's data structure so deeply that the following improvements became apparent:
 - Bulkdata load times improved 3 to 5 times
 - Nodes size scaled down from 12 - 24 to 4 - 8
 - True Market Value generation time was cut in half from 100 minutes to 50 minutes
- Respond to the client's needs with professional analysis of the business domain and custom IT architecture to produce the following infrastructure savings:
 - 200 fewer hosting companies utilized
 - \$10,000 to \$12,000 per month saved (based on 75% utilization)
 - Network bandwidth reduction from 6 Gbps to 3 Gbps for vehicle storage
 - Enabled microservices architecture
- Run multiple learning loops/proofs of concept in parallel to pinpoint the ideal solution
- Perform the phased migration of a production platform with no downtime or performance degradation while successfully feeding learnings from the initial phases back to the migration roadmap and process in a way that didn't affect any day-to-day business operations

WEB APP PERFORMANCE

CALCULATORS-WEB

PAGE TYPE	TTFB - COHERENCE	TTFB - MONGODB
Simplified-pricing	178	83
Select	72	15
Calculator	166	58
Financing-model	161	55
Rest	68	10
Pricing-model	97	56

EFS-REST-WEB

PAGE TYPE	TTFB - COHERENCE	TTFB - MONGODB
Configurator data	84	49
Other data	95	36
Region data	69	6
TMV data	82	37
Vehicle data	177	17
Incentive data	77	12

COMMON-REST-WEB

PAGE TYPE	TTFB - COHERENCE	TTFB - MONGODB
Rest calls	74	12

MOBILE-REST-WEB

PAGE TYPE	TTFB - COHERENCE	TTFB - MONGODB
Media	165	24
Trim	157	66
VD	73	67
Detail	141	28
Model	75	10
Reviews	69	9
Api	73	14
Mobile-rest-tco	181	48
Incentives	69	14

QUESTIONS?
CONTACT US AT
SALES@EPAM.COM

 For more information,
PLEASE VISIT EPAM.COM

41 University Drive, Suite 202,
Newtown, PA 18940 USA
P: +1 267 759 9000 | F: +1 267 759 8989

© 1993-2015 EPAM. All Rights Reserved.



CASE STUDY:

EPAM HELPS EDMUNDS MOVE THEIR DATA CLOSER TO THE BUSINESS WITH A NoSQL SOLUTION

EPAM OPEN-SOURCE DATABASE SOLUTIONS

From the initial professional analysis to the final cleanup, EPAM researched, implemented, and customized the right solution for Edmunds, resulting in more effective, improved internal data structures and solution architecture along with a significant reduction in licensing, operational, and infrastructure costs.

The Coherence Migration for Edmunds is a highly specialized EPAM database migration project, and it sets a precedent for all future migrations. As open-source storage solutions like MongoDB become more and more mature, opportunities abound for companies to transform their database management systems, and there's no better partner than EPAM. Fill out an inquiry today so we can begin our own analysis for your business!