GIGASPACES
innovate with confidence

GIGASPACES AUTOMOTIVE

Groupe PSA Modernizes their Mainframe with GigaSpaces Smart Cache, Delivering 15 Times its Mainframe Capacity While Eliminating Regulatory Fines

GROUPE PSA MODERNIZES THEIR MAINFRAME

USE CASE: AUTOMOTIVE

In 2017, the EU replaced the outdated CO2 emission test NEDC with the new WLTP Protocol (Worldwide Harmonised Light Vehicle Test Procedure), basing it on real-driving data rather than on theoretical driving. In parallel Clean Air For Europe Programme (CAFE) was established aiming to constrain vehicle suppliers and manufacturers, to produce less polluting vehicles by assigning them an average CO2 rate per kilometer emitted annually on all sold vehicles. For example, an extra gram of CO2 emission per kilometer might mean a fine of €100 per vehicle, which can result in huge fines.

If a car manufacturer sells a million cars that emit 1 extra gram of CO2 per kilometer, this can translate to an annual

€100M FINE

PSA GROUPE

About Groupe PSA

Groupe PSA is the second largest car manufacturer in Europe. Present in 160 countries, with 16 production sites across the world, PSA sold 3.5 million vehicles worldwide in 2019.

Business Challenge

Regulations require car manufacturers to produce less polluting vehicles. They are also required to calculate the CO2 emission of each configured car. Failing to calculate and display the CO2 emission can lead to heavy fines.

"The databases we compared GigaSpaces to were much more costly, and were less efficient in terms of performance."

Frédéric Warin, Managing Enterprise Architect at Capgemini for the PSA Project

GROUPE PSA MODERNIZES THEIR MAINFRAME

USE CASE: AUTOMOTIVE

Technical Challenge

Customers, dealership staff and partners use an online web configurator to build and price cars. The mainframe platform that hosts the PSA car configurator had a maximum capacity of 200 requests per second, while PSA estimated actual demand would hit 3000 requests per second. In addition, the mainframe response time was very slow, inadequate to support an online web experience.

Why Did Redis Cache Fail?

The first attempt to solve the challenge using Redis cache failed. Redis is designed for primary key access with simple key lookups, without secondary indexing. Working around this limitation requires duplicating the data multiple times, increasing the footprint and impacting performance. The number of possible configurations reached billions, causing the Redis cache to be inefficient, and not a viable option.

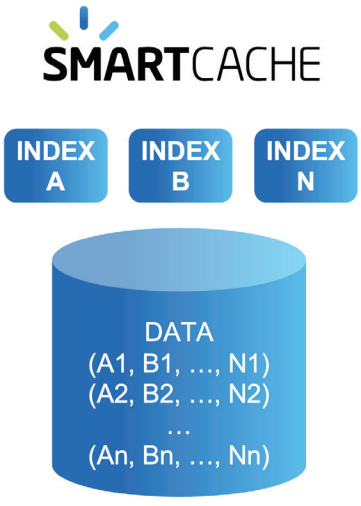

A, B, …, N ALL indexed >>

FAST

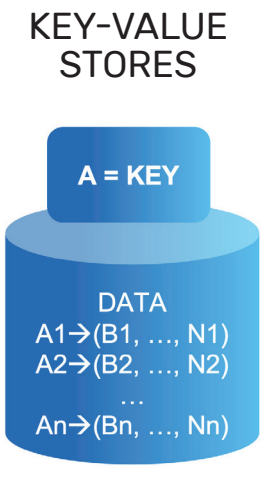

A in the key >>

FAST

B, C, … N are not >>

SLOW!

GROUPE PSA MODERNIZES THEIR MAINFRAME

USE CASE: AUTOMOTIVE

Solution

GigaSpaces Smart Cache was implemented to resolve the bottleneck problem by offloading mainframe requests to its core in-memory data grid engine. A mask was applied to the configuration options, eliminating many of which were irrelevant to the CO2 emission calculation. As a result, the number of configurations dropped significantly, allowing the majority of them to reside in the in-memory multi-model data fabric. This dramatically reduced the number of calculations that needed to be sent to the mainframe. Dynamic multi-criteria queries allowed for superfast access to the data, without duplicating data per each possible combination of criteria.

For high availability, the data is automatically backed up to a database so in case of a server failure, the cache can be restored automatically from the database.

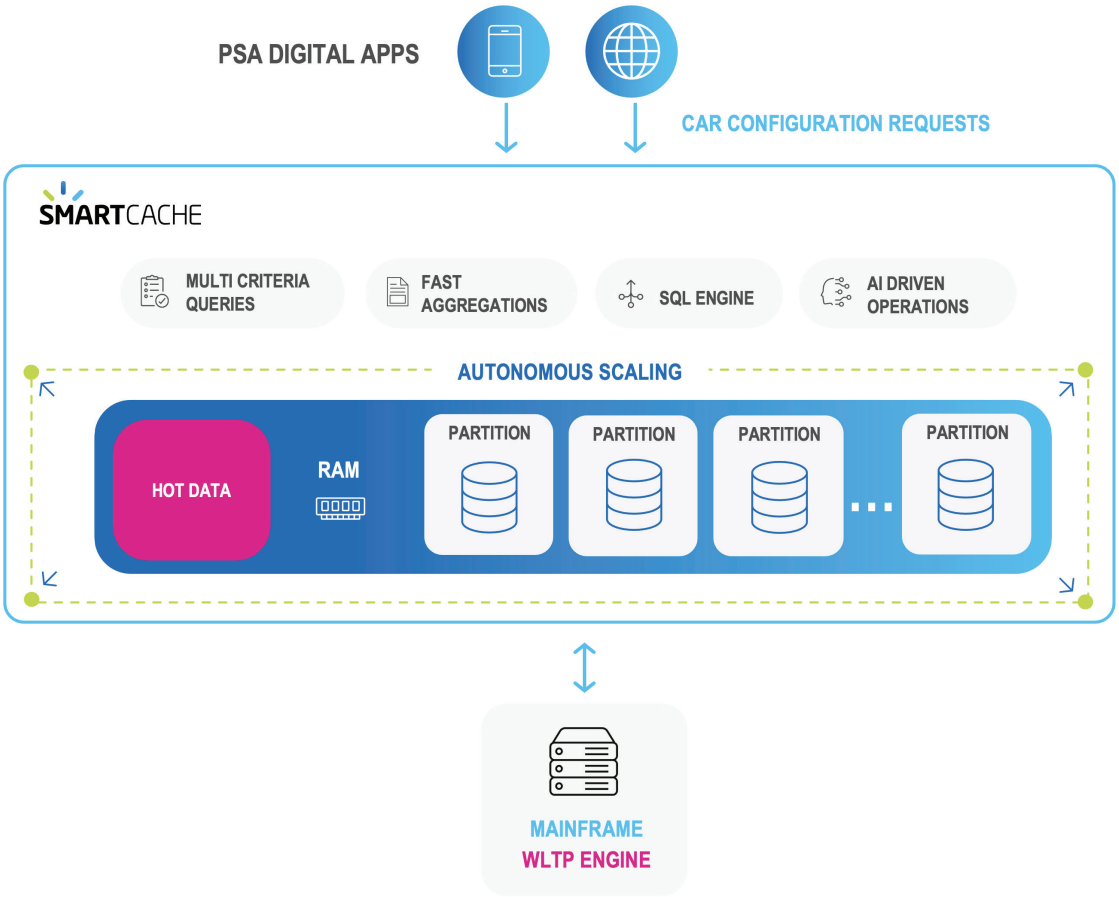

Figure: High Level Solution Architecture

GROUPE PSA MODERNIZES THEIR MAINFRAME

USE CASE: AUTOMOTIVE

Results

Over 95% requests for CO2 calculations are served by InsightEdge Smart Cache

Smart Cache response time ranges between **15-19 milliseconds** vs. 200-300 Milliseconds

Infrastructure footprint was reduced by a factor of **6X** while workload scale was increased **by 20X**

Calculator capacity **increased 15x** without upgrading the mainframe platform

Solution was developed & deployed within **12 weeks** from kickoff to go-live

ABOUT GIGASPACES

GigaSpaces is redefining in-memory technology to drive enterprise digital transformation with unparalleled speed, performance and scale. Never before has it been this simple to accelerate and scale real-time applications, analytics and operational BI on any data, at any load, across any environment. Hundreds of Tier-1 and Fortune-listed organizations and OEMs across financial services, retail, transportation, telecom, healthcare rely on GigaSpaces to optimize their business operations, enhance customer experience and comply with regulations.

"GigaSpaces' InsightEdge makes you faster and smarter at the speed and scale of business with strengths in transactions, platform, multi-model, customer adoption, development tools, extensibility, and customer support."

The Forrester Wave™:
Translytical Data Platforms
Q4 2019

300+
Direct Customers

50+/500+
Fortune/Organizations

5,000+
Large installations in production (OEM)

25+
ISVs

BENEFITS

EXTREME SPEED

Ultra-low latency, high-throughput transaction, stream and analytical processing. Co-location of data, applications and analytics eliminates the need to move data to act on time-sensitive data at milli and microsecond performance. Event-driven analytics and ML runs simultaneously on live, hot and historical data on data lake.

ANY DATA

Ingest, process and analyze from any data source including structured, unstructured or semi-structured. Applications and analytics seamlessly access documents, images, tables, objects geo-location and more.

EXTREME SCALE

Elastic linear scaling and handling of peak events without performance degradation is supported by a distributed in-memory scale-out, shared nothing architecture across any environment.

MISSION CRITICAL AVAILABILITY

Mature battle-tested platform for mission critical businesses. Highly available with up to 5 nines reliability and auto healing. Geo-redundancy and fast data replication for disaster recovery.