

MongoDB for Automotive Manufacturing

Driving innovation from
factory to finish line



Table of Contents

| | |
|---|----|
| The State of the Automotive Industry | 3 |
| Overview of MongoDB | 4 |
| Use Cases for MongoDB in Automotive Manufacturing | 5 |
| Smart Supply Chain | 5 |
| Track and Trace | 5 |
| Inventory Management and Optimization | 7 |
| Logistics Optimization | 8 |
| Planning and Forecasting | 9 |
| Smart Factory | 11 |
| Connected Equipment | 11 |
| Connected Workforce | 13 |
| IIoT Enabled Operations (Next Generation Manufacturing Execution System) | 15 |
| Manufacturing Intelligence | 18 |
| Edge Analytics | 18 |
| Predictive Maintenance | 19 |
| Adaptive Quality Assurance | 20 |
| Virtual Factories | 21 |
| Smart Commerce | 23 |
| Customer 360 | 23 |
| Product Catalogs | 25 |
| Inventory Management | 26 |
| Smart Products and Services | 27 |
| Connected Mobility and Digital Customer Experience | 27 |
| Conclusion | 29 |
| About the Author | 29 |



The State of the Automotive Industry

Modern automotive manufacturers need to innovate quickly and effectively to stay competitive in today's marketplace. One of the most important tools for enabling this kind of innovation is powerful and easy-to-use real-time data analytics tools that allow companies to quickly access industry-level data, enabling them to make informed decisions.

No longer are companies limited to retrospective analysis of data, instead they are striving to make real-time decisions based on the most accurate and up-to-date information possible. The ability to explore large amounts of data within milliseconds not only accelerates the pace of innovation, but also ensures that companies can develop their products with the most precise and accurate information when making critical decisions on future business actions.

The automotive value chain is complex and involves a lot of moving parts. In the product design and development phase, the challenge is to manage data from multiple sources and formats including CAD files, simulations, product test results and customer feedback. Many legacy systems are incorporated at this stage to handle data movement and aggregation. Unstructured data needs to be contextualized and consumed via various applications. Not having a streamlined data infrastructure results in siloed data, redundant work, and slow collaboration between departments. In the second phase of planning and procurement, the major challenge is to effectively coordinate and collaborate with suppliers. Achieving real time visibility on raw material shipment and inventory tracking can result in optimized and efficient production but this is a challenge because suppliers may use different data formats or have varying levels of data quality resulting in data inconsistencies and inaccuracies.

During manufacturing, there is a need to manage high volumes of time series data generated by machines, robots, and people working on the shopfloor. These days, machines are being produced with embedded sensors and the

cost of external sensorization is dropping. This presents a unique opportunity and a challenge for automotive manufacturers around the world. The challenge is not just to work with multiple data sources but to also simplify the data aggregation process in a way that it can be deployed easily for the company, thus enabling data driven insights to improve daily production metrics. Quality inspection systems need to be integrated with the fabrication and assembly lines in order to keep the production parameters optimized at all times. Once the vehicle is manufactured, in the final stages of the value chain, the challenge is to manage the supply chain, logistics, and customer relationships. Real-time data is required to track the movement of vehicles, monitor inventory levels, and forecast demand. This requires connectivity from the vehicle back to the central data platform which is monitoring the health and performance of the vehicle generating actionable insights for the manufacturer to improve their product and keep their customers satisfied. Tracking customer preferences, feedback, and behavior can be challenging due to data being locally stored at dealerships and not flowing to OEM. This can lead to suboptimal product offerings or ineffective marketing campaigns.

In summary, there are data related challenges at each stage of the automotive value chain which leads to inefficiencies, quality issues, and delivery delays throughout the value chain. By addressing these challenges through the use of flexible and versatile data platforms, advanced analytics, and automation, manufacturers can innovate faster, improve production efficiency and performance and stay competitive in a rapidly evolving industry.



Overview of MongoDB

MongoDB helps automotive manufacturers drive innovation. IoT technologies, like condition monitoring sensors, allow automotive manufacturers to make informed decisions faster, providing a competitive edge over competitors. Additionally, real-time access to data through MongoDB enables much better collaboration between teams, speeding up the innovation process. With MongoDB, automotive manufacturers can capture bigger and better opportunities for creating groundbreaking products and services, thus helping them stay on the cutting-edge of their industry.

The tabular row-and-column data model used by relational databases bears little resemblance to how data is represented in application code. In modern programming languages, the entity you want to store in the database (i.e. customer, product, order, trade, log message, sensor reading) is represented as a complete object, with all related attributes contained in a single data structure. In contrast to the tabular model, the document data model presents a much more intuitive and natural way to describe data. This is because documents are consistent with the way developers think and code. Rather than splitting data apart and flattening it out across multiple tables, documents are closely aligned to the structure of objects in the programming language. Documents are single and contain data structures with related data embedded as subdocuments and arrays. In more advanced document databases like MongoDB, each element can be individually indexed and updated, no matter how deeply nested it is within the document. As a result, the document model approach is simpler and faster for developers to model how data in the application will map to

data stored in the database. It also significantly reduces the barrier-to-entry for new developers who begin working on a project.

MongoDB Atlas is a multi-cloud developer data platform. At its core is our fully managed cloud database for modern applications. MongoDB Atlas is the best way to run MongoDB and the document model is the fastest way to innovate because documents map directly to objects in your code. MongoDB's unified query API is the most natural way to work with data in any form. MongoDB Atlas extends MongoDB's flexibility and ease of use to build full-text search, real-time analytics, and event-driven experiences.

The MongoDB Atlas database is available in 100+ regions across AWS, Google Cloud, and Azure. This flexibility also allows you to take advantage of multi-cloud and multi-region deployments, allowing you to target the providers and regions that best serve your users. This best-in-class automation and proven practices guarantee availability, scalability, and compliance with the most demanding data security and privacy standards.

“Previously, launching new services for a customer would take two weeks. With MongoDB, we are done in 30 minutes: our field engineers can do all of this in just a matter of clicks on a console with absolutely zero development effort. We reduced more than 90 percent of requests coming to us just by giving them this tool, which saves us all time and money.”

—RANJAN MOSES, WEKAN'S CHIEF TECHNOLOGY OFFICER,
ON USING MONGODB FOR [RIDEKLEEN AT COX AUTOMOTIVE](#)



Use Cases for MongoDB in Automotive Manufacturing

MongoDB enables a multitude of use cases across the automotive value chain as listed in Figure 1.

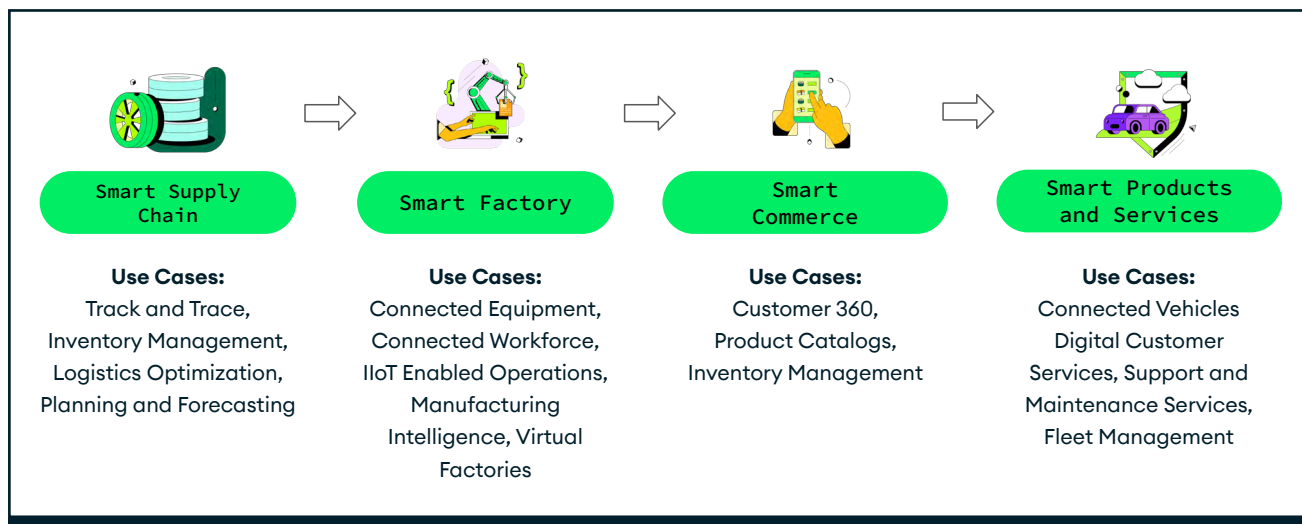


Figure 1. MongoDB enabled use cases for automotive manufacturing value chain

Smart Supply Chain

Effective management of supply chains is critical to ensure the timely delivery of parts and materials required for vehicle production. While numerous business processes have to work together to manage the supply chain, here we will cover four key use cases that MongoDB can enable for automotive manufacturers.

Track and Trace

Asset tracking in automotive manufacturing is no longer a mere location tracking tool. When combined with other disruptive technologies such as cloud computing, advanced analytics, and NB-IoT (Narrowband IoT), businesses can extract valuable insights from the data to optimize operational efficiencies.

With the evolution of technologies such as BLE (Bluetooth Low Energy), LoRaWAN (Long Range Wide Area Network), and NB-IoT (Narrowband IoT), the reliability and performance of asset tracking solutions have greatly enhanced. Low power tags connected with a positioning engine enable automatic and precise tracking of assets across

the manufacturing shop floor. Such tags are mobile with long battery life and can be attached to any movable equipment, like forklifts, trolleys, and dollies.

Typically, companies need to add context to location data coming in from asset management systems, which means they need an operations data layer to aggregate data from different sources. In most cases, they have to incorporate an operational context such as the status of the asset (operational, under maintenance or out of service), operational zone, or specific area within the factory to provide a more comprehensive understanding of the asset's location in relation to its operational state. This helps in optimizing workflows, identifying bottlenecks, and

streamlining operations. MongoDB enables this context addition to asset track and trace through a wide variety of features. MongoDB Atlas sits in the middle as a central operational data hub collecting data from various sources including the localization sensors and the inventory data of car parts from the supplier warehouse. All this data can be modeled effectively with the Document Data Model and stored in an optimized manner such as time series collections. MongoDB documents are inherently flexible while allowing data governance when required. In the factory, Kafka Connector can be used to stream sensor data to Atlas whereas real-time data changes in other data stores can be pushed to MongoDB

through a Change Data Capture (CDC) technique. MongoDB Change Streams can be used to write data back to other enterprise software systems such as Order and Inventory Management. As the car parts from suppliers go out for delivery, the location data can be stored on a mobile device using Realm SDK and synced to Atlas through the use of App Services and Device Sync.

Finally using Atlas Search and Atlas Vector Search, the users of the system can retrieve status information about assets locations by performing search based on matching keywords/fuzzy matching or by performing a similarity search on the stored vector embeddings.

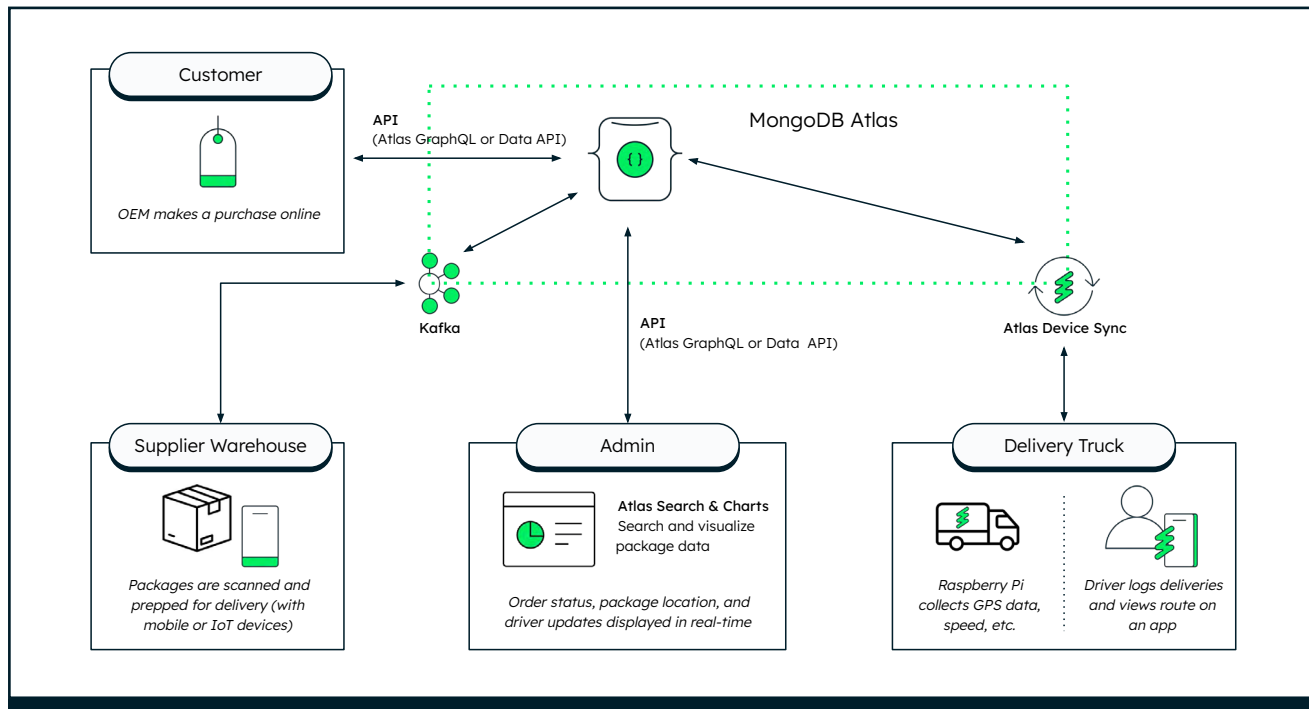


Figure 2. MongoDB as the central data store for Asset Track and Trace use case

Inventory Management and Optimization

Modern automotive supply chains are complex systems that interconnect the globe. Efficient supply chains control costs and ensure on time delivery in full to customers. Inventory management is a key component in achieving a well functioning supply chain. Higher inventory levels allow for car parts suppliers to maintain better customer service levels but those come with higher cost of holding inventory. Thus each automotive part supplier has to ensure optimal balance in inventory levels to maximize profitability and maintain market competitiveness. They also need to efficiently coordinate across the supply chain to minimize the bullwhip effect which can result in amplified demand variability upstream.

For automotive manufacturers, to manage and optimize inventory levels, the first step is to maintain an accurate and real-time view of inventory levels across multiple plants, warehouses, and suppliers. This is absolutely essential as without having visibility on the current inventory levels, it is impossible to optimize. The second step is to reduce inventory carrying costs while still ensuring that the required parts and materials are available to ship out when needed. Finally, the data from multiple customers needs

to be aggregated and analyzed despite being in different formats, each with its own unique lead times and order quantities.

MongoDB Atlas provides a flexible, scalable, and highly available developer data platform for managing inventory data. The document data model can handle complex inventory structures and hierarchies, making it easy to manage inventory across multiple plants and suppliers. At the warehouse, the inventory can be scanned using a mobile device. This data is persisted in Realm and synced with Atlas using Device Sync. Realm and Device Sync provide an offline-first seamless mobile experience for inventory tracking, making sure that inventory data is always accurate in Atlas. Once data is in Atlas, it can serve as the central repository for all inventory-related data which includes stock, supplier, and customer information, bill of materials and production line data. MongoDB Application and API services enable connectivity and streamline collaboration between different departments and systems, such as supply chain management, production planning and sales. Atlas Search can be used to perform full text search on inventory data stored in Atlas, reducing time to find the required information. This functionality can be extended using Atlas Vector Search to enable a chat-bot like experience which can be used to ask questions about inventory status.

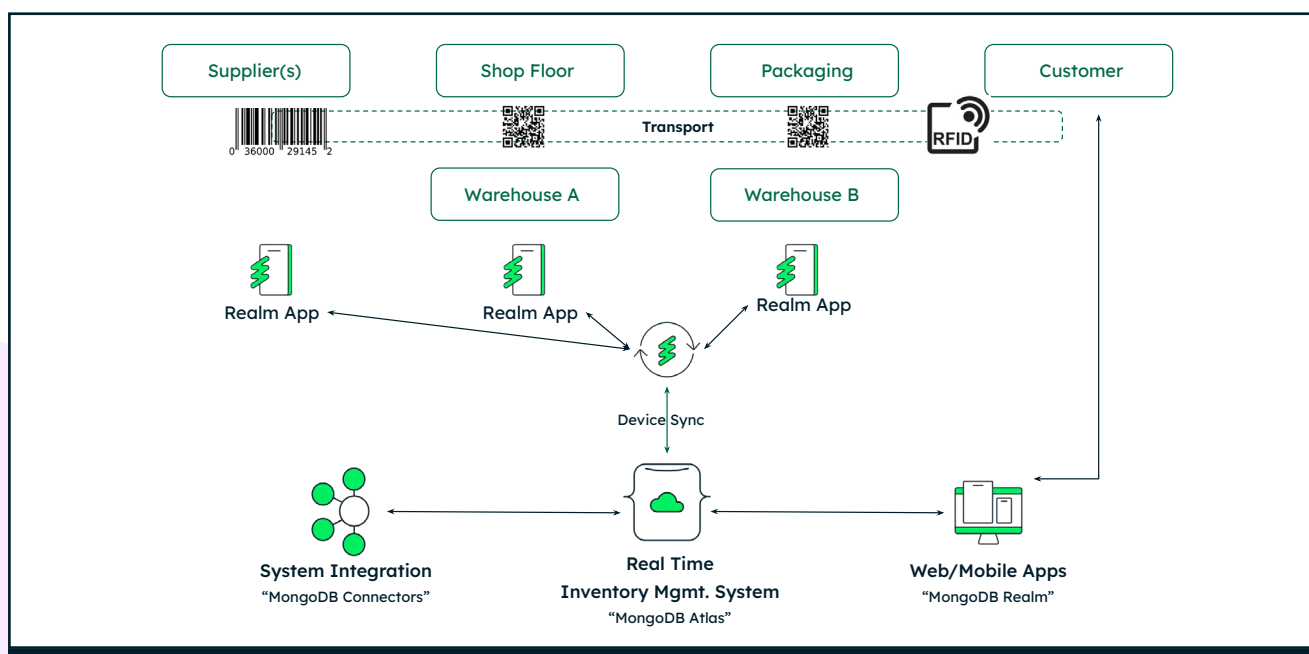


Figure 3. MongoDB Atlas and Realm enabling real time inventory management



Logistics Optimization

The demand for highly personalized and customized vehicles is continuously growing. Thus, both inbound and outbound logistics management systems have to adapt to this changing environment. Logistics optimization makes sure that customer vehicle orders are delivered on time in full (OTIF). This results in higher revenue and profitability for the automotive manufacturers. At the storage, warehousing, and material handling stage, it is important for manufacturers to have a healthy stream of products coming to build inventory. At the transport stage, the transportation management system (TMS) needs to enable interactions between an order management system and the warehouse or distribution center so customer orders are picked up and delivered on time using a mix of air, road, and rail networks. When the product is on the road on its way to delivery, companies need to use GPS technology to accurately locate vehicles, monitor freight movement, negotiate with last mile logistics companies, and consolidate shipments.

The first step towards optimizing logistics is to enable real-time visibility from production to delivery. Once the inventory, transportation, and delivery data is available, data analytics can be performed to enable route optimization, manage warehouse space efficiently, and improve overall supply chain visibility and responsiveness.

Logistics optimization relies heavily on accurate and real-time data from production all the way

to delivery. This data includes inventory levels, transportation routes, customer orders and demand forecasts. For automotive manufacturers, obtaining reliable data can be a challenge, as it may be scattered across different systems and not easily accessible or integrated. Such inaccurate and incomplete data hinders the effectiveness of optimization efforts.

MongoDB Atlas provides a flexible, scalable, and highly available developer data platform to help automotive manufacturers optimize their logistics systems. The flexible document data model allows for modeling and storage of diverse logistics data from Warehouse Management System, Transport Management System and Order Management System and others. As the data grows, MongoDB's horizontal scalability (sharding) allows for manufacturers to handle high data volumes and increasing workloads. With regional sharding and multi-cloud deployment sharding, the data storage can be optimized and regional storage enabled for regulatory compliance. Route optimization applications involve geospatial analysis such as route planning and vehicle tracking. In MongoDB, you can store geospatial data as GeoJSON objects and the database provides powerful geospatial indexing and querying features, facilitating location based logistics optimization. Finally the Kafka connector enables easy integration with third party TMS, WMS and other systems to bring all this disparate data into MongoDB.

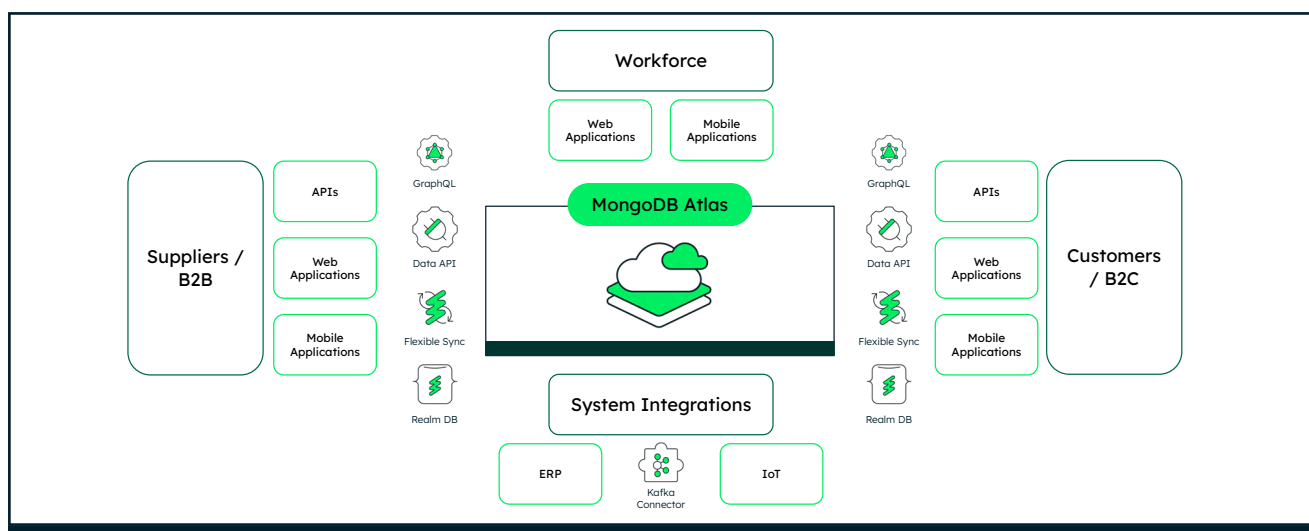


Figure 4. MongoDB Atlas as the logistics data layer



Planning and Forecasting

Many automotive manufacturers operate in a Just In Time (JIT) environment and do not carry safety stock. Because of that, they expect their suppliers to provide the parts in a timely manner to hit production targets. There is a need to have accurate and timely information flow upwards from the customer order to the suppliers and for this information to be of any value, it has to be provided early enough so that suppliers can account for adjusting production and order quantities as the demand changes. Additionally, this information needs to represent future demand as well so that suppliers and other supply chain partners can react accordingly. If this information

is not accurate or real-time, the bullwhip effect will occur resulting in inaccurate inventory forecasts, unoptimized shipping plan and over/under production.

For both automotive suppliers and OEMs, supply planning and demand forecasting requires processing and analyzing diverse data including historical sales data, market trends, promotions, customer behavior, competitive information, and other external factors. Analyzing this data can be extraordinarily challenging if the data that holds the answers is scattered among different incompatible systems, formats and processes. Essentially this turns into a huge data engineering problem.

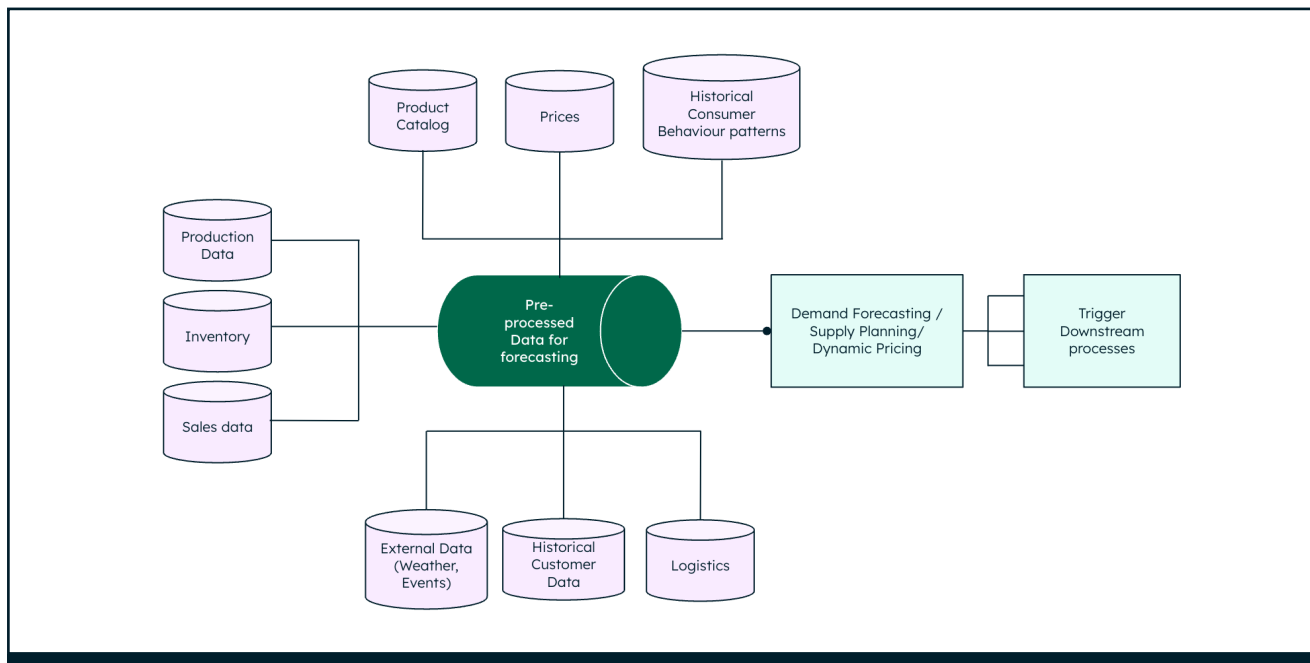


Figure 5. Supply Chain Data Processing - A Data Engineering Challenge

Consolidating all this data in a central location to make it available for supply chain analytics applications is typically done by batch oriented ingestion methods. MongoDB with its scalable and flexible document data model can handle the complexity and variability of this data, allowing for real time ingestion, efficient storage, retrieval and analysis of diverse data sets. MongoDB Atlas allows for schema validation during ingestion to ensure there are no unintended schema changes or improper data types being stored.

MongoDB Aggregation framework simplifies data pre-processing. It allows filtering, grouping, joining, and sorting data, while cleansing, and transforming it within the database as well as creation of materialized views for low latency queries.

After preprocessing, automotive manufacturers and their suppliers can use the MongoDB Atlas ecosystem integrations, such as Kafka, Spark and SQL interface to connect to popular analytics

Use Case

Longbow Advantage delivers substantial business results by enabling clients to optimize their supply chains. Its flagship warehouse visibility platform, Rebus, combines real-time performance reporting with end-to-end warehouse visibility and intelligent labor management.

[Read the full story.](#)

tools and frameworks, enabling their supply chain analytics team to leverage their existing data analytics and visualization tools to derive insights from their data. Figure 6 below shows how MongoDB Atlas can sit in between Ecommerce systems and ML platform (Databricks) for storing raw ecommerce data and doing preprocessing for ML based dynamic pricing for car accessories.

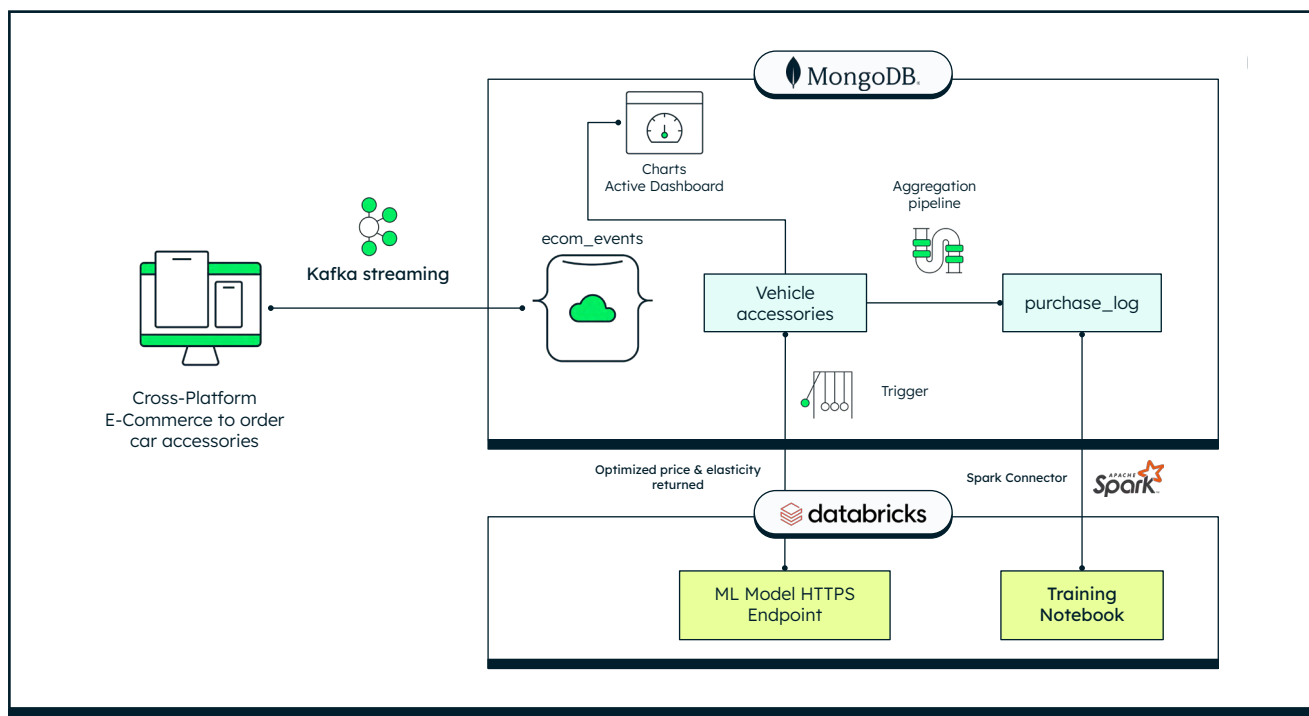


Figure 6. Dynamic Pricing with MongoDB Atlas

Smart Factory

By implementing the Industry 4.0 concepts, automotive manufacturers can enable increased productivity, efficiency, and production cost savings. MongoDB helps our manufacturing clients to optimize production processes, reduce downtime, improve quality control and enhance worker safety. Below we are going to cover five key use cases that MongoDB can enable for automotive manufacturers.

Connected Equipment

Industry 4.0 generally comprises many complex components and has broad applications in automotive manufacturing. The very first challenge faced by manufacturing companies when embarking on the I4.0 journey is to sensor and connect their manufacturing equipment in order to collect, store and analyze data for information and insights.

IIoT data from equipment, products, and services can be used in novel ways to identify unique business opportunities and enable digital ways of working. Through the use of data-driven decision making, operational efficiency can be gained for production and supply chain management. Gone

are the days when manufacturing companies were only relying on two or three sources of production data to improve their Overall Equipment Effectiveness (OEE). The traditional method of manually calculating OEE using spreadsheets seems outdated in the digitalization age.

These days, machines are being produced with embedded sensors and the cost of external sensorization is dropping. This presents a unique opportunity and a challenge for manufacturers around the world. The challenge is not just to work with multiple data sources but to also simplify the data aggregation process in a way that it can be deployed easily for the company, thus improving developer efficiency along the way. Distributed data sets sitting in silo data stores are difficult to access and keep in sync.

MongoDB Atlas is well positioned to solve such challenges because it is built around the most intuitive way to model data - the document data model. It provides a unified query interface for a broad range of workload types (time series, operational, real-time analytical, etc.) and supports a wide range of modern IIoT application types as they grow and evolve. MongoDB is a highly scalable document based database that can handle large amounts of structured, semi-structured and unstructured data. Native time series collections are available that help with storing large amounts of data generated by IIoT enabled equipment in a highly compressed manner. The database flexibility makes it well-suited for handling the different data structures needed to store sensor data, ML models and prediction results, all in one database. This removes the need for maintaining separate databases for each type of data reducing IT sprawl in automotive manufacturing

Use Case

Industry 4.0 has arrived and Toyota Material Handling Europe is building both the machines and the data platforms that will make this historic shift in manufacturing a success. Toyota is famous for the quality of its vehicles. Now that same standard is being applied to how its development teams create Internet of Things software that will enable the move to smart, autonomous and safer factories of the future.

“At Toyota, we aim to be number one in industry 4.0,” explained Filip Dadgar, Principal System Architect and IT-Manager at Toyota Material Handling Europe.

[Check out the story.](#)



organizations. Finally, analysts can use the Atlas SQL interface to access MongoDB data from SQL based tools. This allows them to work with rich, multi-structured documents without defining a schema or flattening data. In a connected equipment setting, this ability can be useful to generate reports for equipment failures over a period of time and comparison between different equipment failures types. Data can be blended

from MongoDB along with other sources of data to provide a 360 degree view of production operations. The power of MongoDB can be utilized on the shop floor level and if needed continuously be replicated to MongoDB Atlas on Azure, AWS, GCP or a self managed MongoDB cluster anywhere which can be synced with Atlas using cluster-to-cluster sync functionality.

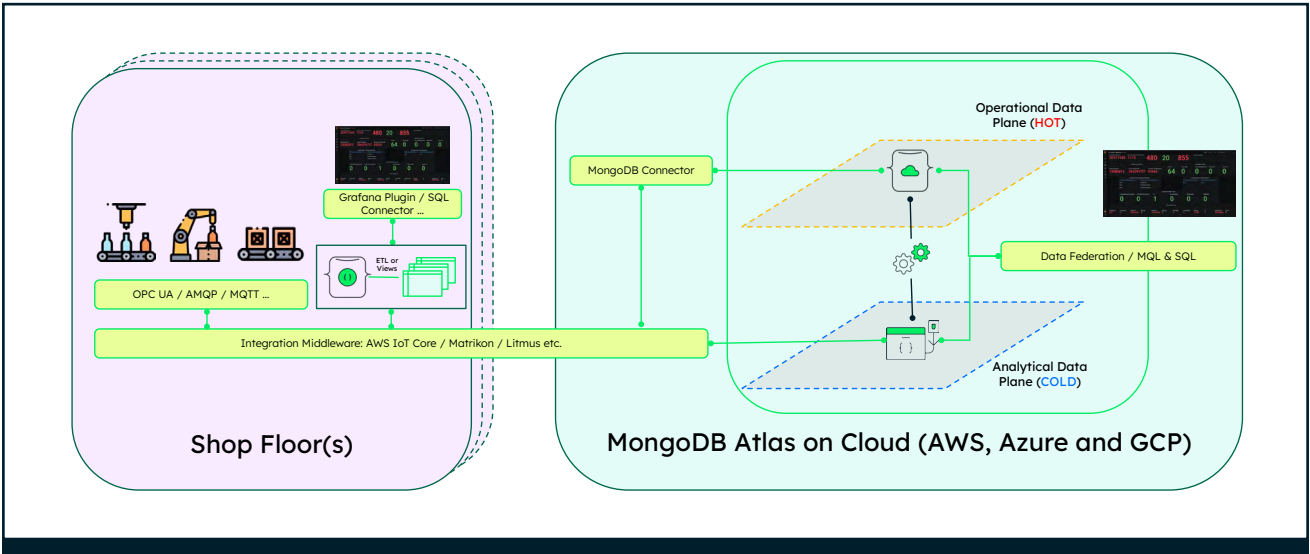


Figure 7a. MongoDB Atlas as the Data layer from shop floor to cloud

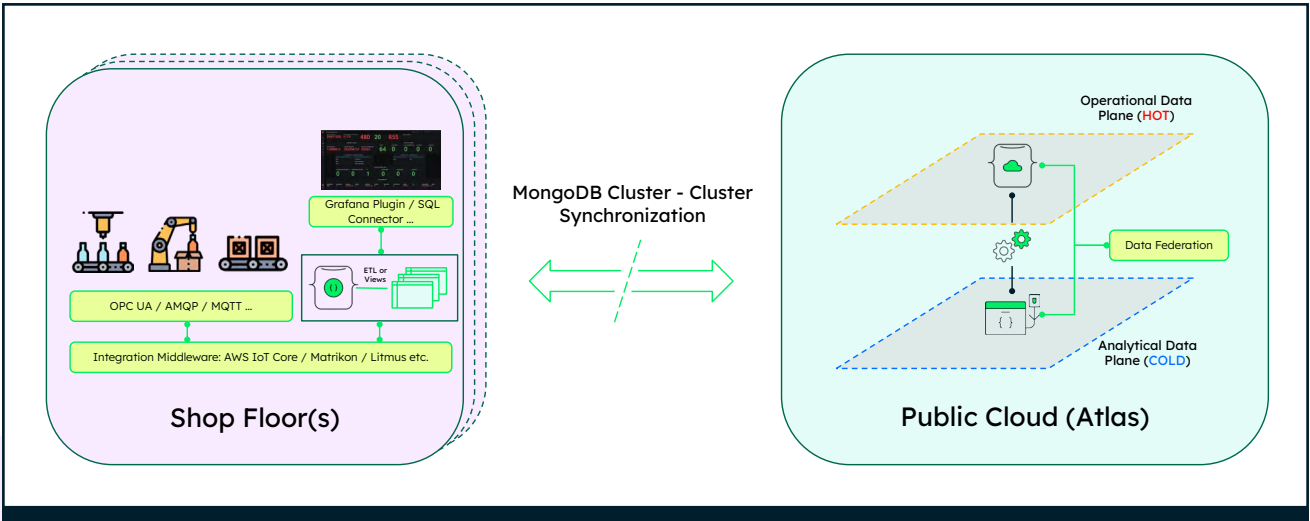


Figure 7b. MongoDB Atlas as the Data layer from shop floor to cloud (using Cluster-to-Cluster Sync)



Connected Workforce

Automotive manufacturers need digital tools and technologies to help their workers collaborate and stay connected across regions, geographies and functions. These tools help to get workers

onboarded, collaborate in real time for planning, drive operations and maintenance activities and to tap into the collective knowledge of the enterprise for solving complex challenges. The table below shows examples of connected workforce mobile applications.

| Automotive Business Unit / Department | Examples of Mobile Based Applications | Value Unlocked |
|---------------------------------------|---|--|
| Product Research and Development | <ul style="list-style-type: none">• Review design iterations and mockups• Collaborate between Design and Development | <ul style="list-style-type: none">• Increased daily interactions between functions• Decreased redesigns and cost associated• Increased revenue due to less time to market |
| Incoming Order Processing | <ul style="list-style-type: none">• Receiving and entering incoming customer orders• Communicating delivery schedules and changes to orders | <ul style="list-style-type: none">• Increased daily interactions between functions• Increased revenue due to fast order capture and customer satisfaction |
| Inventory Operations | <ul style="list-style-type: none">• Receiving and entering inventory information• Prioritizing products for shipping and initiating shipping workflow• Monitoring inventory levels | <ul style="list-style-type: none">• Freed working capital• Revenue generated by increased inventory turnover |
| Maintenance Operations | <ul style="list-style-type: none">• Routine (Daily/Weekly) maintenance checks on production equipment• Troubleshooting equipment failures• Responding to maintenance alerts and digitally collaborating with experts to fix the issue | <ul style="list-style-type: none">• Decreased maintenance costs• Decreased Mean Time to Repair (MTTR)• Decreased travel cost of experts• Increased daily interactions between functions |



All these examples sound great but building intuitive mobile apps from the ground up presents a lot of challenges for already strained IT teams. Legacy architectures in typical manufacturing tech stacks are often made up of relational and niche NoSQL databases, and additional analytical platforms – all resulting in siloed data, slow data processing and unnecessary complexity. For example, in a mobile application that helps with troubleshooting equipment failures, the data from Central Maintenance Management System (CMMS) and IIoT platform needs to be combined together and kept in sync with the mobile application database. This requires time-consuming ETL maneuvers to bring data together into a single view for the mobile application. It is also hard to ensure data synchronization between a worker's app and the backend database when they're moving in and out of connectivity (when workers walk to the back of a warehouse for instance where WiFi coverage might not be available). It's even harder with a sprawling data architecture to account for. Finally the added complexity of managing multiple databases and connections between mobile devices and the cloud backend slows down the innovation speed at an organization.

MongoDB's Developer Data Platform provides functionality to create an application backend that is synchronized with the database, alongside several features so that automotive companies can push out workforce applications quickly and allow developers to focus on differentiating code when building a mobile app.

MongoDB solutions for mobile applications development consist of three things:

- **Realm Mobile Database** a lightweight embedded database that sits on the device to make sure the app is always available and fast, regardless of network connections.
- **Atlas Device Sync** allows you to sync data between your Realm Mobile Database and the cloud. Sync keeps the data updated across users, devices, and the backend.
- **MongoDB Atlas** – Our fully managed cloud database – on the backend, so developers can stop worrying about scale as app usage changes. This also opens the teams up to leverage the full extent of the Atlas platform – such as Atlas Search, Triggers, or Charts, just to name a few.

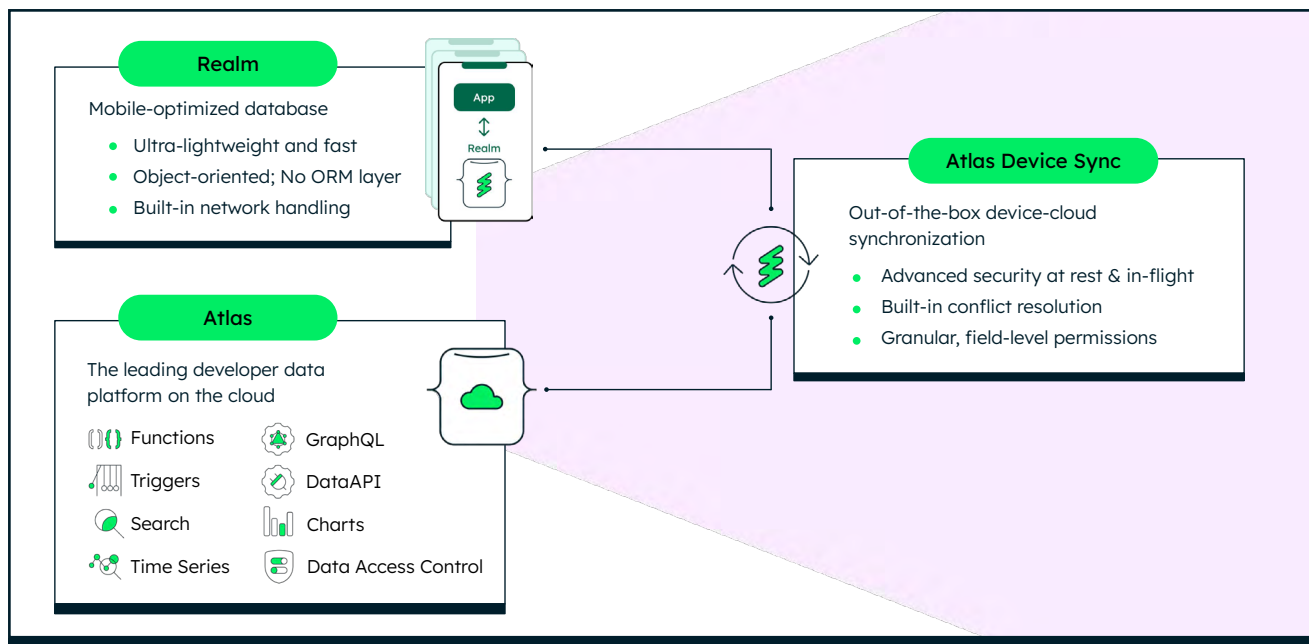


Figure 8. MongoDB is the easiest way to build and scale mobile apps



Because Realm and Atlas Device Sync are built to work cross-platform, automotive manufacturers can have one data layer that can be consistent and reused across all workforce devices, ensuring that the data is consolidated in a single platform and all teams have visibility into the valuable data collected from devices.

At the end of the day, MongoDB is committed to the developer experience and making building technology simple. When it comes to mobile apps, we provide a reliable real time mobile application experience, a stellar developer experience and a unified platform on Atlas to simplify maintenance and scalability.

IIoT Enabled Operations (Next Generation Manufacturing Execution System)

For the last three decades, Manufacturing Execution Systems (MES) have been an essential industrial software for production execution and shop-floor automation. However, the advent of the Industrial Internet of Things (IIoT) has created a demand for a next-generation MES that can process large amounts of time series sensor data generated by IIoT devices and merge it with production execution data. Integrating IIoT technologies with legacy MES software has proven challenging for automotive manufacturers due to the monolithic design, inflexible data model, and poor horizontal scalability of relational databases used to store MES data. The relational databases used in traditional MES systems struggle to handle the massive amounts of IIoT data and product and process varieties that are common in today's

automotive plants. Each generation of connected machines comes with new sensors, bringing greater data volume and variety along with it – as well as potential new business insights. The data generated by these equipment tends to be large in both volume and frequency, placing a strain on the underlying relational database. This data is also time series forcing manufacturers to invest in a specialized time series database just to store IIoT data. This leads to the expense and complexity of yet another data silo, slow and fragile data integration with production operation data in MES, and archiving of aged data.

MongoDB addresses these challenges, allowing IIoT devices and MES to work collaboratively and synchronously, meeting the requirements of a next Generation MES: scalability, flexibility, and interoperability. The following table lists out the trends and critical capabilities required by a next generation MES and how MongoDB enables those capabilities.

Use Case

Bosch is in the business of collecting IoT data and making it actionable, helping engineers and developers to solve problems, design better solutions, and innovate at scale. Learn how Bosch leverages MongoDB to power its IoT applications.

[Read the story here.](#)



| Critical Capabilities of Next Generation MES | Challenges With Relational Databases for MES | How MongoDB Addresses These Challenges |
|---|--|---|
| Flexible production workflow design and deployment | Schema redesign to handle complicated processes, resulting in time and productivity loss. | Data modeling using dynamic documents enables varying object representations in a single collection and removes abstraction. |
| Real time defect detection and quality management with closed loop feedback | Multiple ETL jobs to connect and integrate with brownfield quality inspection databases resulting in increased complexity. | <p>Freedom to run anywhere means you can run MongoDB close to the equipment all the while automatically syncing select collections with Atlas using Cluster-to-Cluster Sync.</p> <p>App Services and Triggers to send events and notifications with shop floor applications and equipment.</p> <p>Atlas Search to retrieve quality issues and logs with ease.</p> |
| Time series IIoT data management | Multiple IIoT sensors and data sources sending data in different formats putting a strain on underlying RDBMS resulting in high latency leading to bad user experience. | MongoDB's time series collections automatically store time series data in a highly optimized and compressed format, reducing customer storage footprint, as well as achieving greater query performance at scale. |
| Offline first mobile applications for workforce | Mobile app databases need to be kept in sync with the central MES database to update the track and trace records resulting in more time spent on keeping applications running and in sync. | <p>Using Realm SDK and Device Sync, mobile apps stay in sync with the central MES Atlas database - all in real time.</p> <p>App UI updates immediately in response to shop floor alerts and notifications and works regardless of network connectivity issues.</p> |



| Critical Capabilities of Next Generation MES | Challenges With Relational Databases for MES | How MongoDB Addresses These Challenges |
|---|--|--|
| Single View analytics and metrics reporting | <p>The metrics reporting module needs to correlate MES data with IIoT sensors/gateways data to present a single view of factory performance to the stakeholders.</p> <p>There is time spent on correlating all this data using complicated SQL scripts at the expense of value added activities.</p> | <p>Aggregation framework gives flexibility to get value out of the data without added complexity and fragility of SQL.</p> <p>For historical data reporting, Atlas Data Federation allows natively querying and combining data across different data sources such as MongoDB Atlas and AWS S3 without complex integrations.</p> <p>Atlas Charts can be used create powerful visualizations right on top of MongoDB data without needing any additional software or code.</p> |
| Flexible and scalable brownfield production equipment integration | <p>Many different types of machines and controllers to integrate resulting in tons of custom software on top of standard MES deployment resulting in increased time to market.</p> | <p>The power of MongoDB can be utilized on the shop floor level and if needed continuously be replicated to MongoDB Atlas on Azure, AWS, GCP or a self managed MongoDB cluster anywhere.</p> <p>MongoDB is designed from the ground up to be a distributed database. MongoDB's horizontal, scale-out architecture can support huge volumes of IIoT data.</p> |
| Scalable and flexible Enterprise data integration | <p>Significant data integration effort is required to combine BOM, recipe, material management, inventory status and order data into one data model.</p> <p>Increased complexity of final system dealing with multiple databases and data models.</p> | <p>MongoDB documents can represent data of any structure while each document can contain different fields. Documents map to the objects in code, so they are much more natural to work with. There is no need to decompose data across tables, run expensive joins, or integrate a separate Object Relational Mapping (ORM) layer.</p> <p>Data that is accessed together is stored together so end users get higher performance.</p> |



Manufacturing Intelligence

The digital economy is demanding that manufacturing applications become smarter, driving better customer experiences, surface insights, and take intelligent action directly within the application on live operational data – in real-time. The objective is to always out-innovate the competition. To meet those demands for working with fresh data, we can no longer rely only on moving data out of our operational systems into analytics stores - this adds too much latency, and it separates the application from the insight that is created. And so to overcome these challenges, analytics processing has to be “shifted left” to the source of the data – to the applications themselves. MongoDB calls this shift “Application- Driven Analytics”. And it’s a shift that both developers and analytics teams need to be ready for because it impacts their roles and responsibilities, along with the tools and technologies they are using. MongoDB serves application-driven analytics through a set of platform capabilities and features - from database through data lake, a federated query service and connectors.

Manufacturing intelligence use cases unearth previously unknown patterns in manufacturing processes or workflows. Let’s look at a few use cases:

Edge Analytics

At the automotive shop floor level, there is a huge amount of important data generated by IIoT enabled smart devices and sensors. Processing and analyzing all this data can be a big challenge for automotive companies. To avoid delays and conduct real-time data analytics, to provide crucial insights, to drive productivity and revenue, edge analytics has gained a lot of momentum in recent years. An Industrial IoT Edge platform is purpose-built to aggregate, analyze and report real-time insights to stakeholders. It can provide

a way to centrally manage and control devices, data and applications across the shopfloor. This includes many business benefits such as near real time analysis, easy scalability, cost reduction and improved security.

MongoDB Atlas provides the best-in-class, flexible development platform that delivers a consistent and painless development experience across the data layer – from data center, cloud, and to the edge.

Atlas Device Sync and Realm: Atlas Device Sync provides out-of-the-box synchronization between devices and the cloud, network handling and retry logic, conflict resolution, and more, paired with Realm, an embedded database that stores data on-device for reliable, responsible experiences regardless of network connection.

Native Time-Series Platform: MongoDB supports native time series collections with hands-free schema optimization supporting high storage efficiency and low latency queries. This helps in collecting time series data from machine sensors and controllers and storing them in a highly optimized manner on the edge.

Aggregation Framework: By using the built-in aggregation framework in MongoDB, users are able to do real time analytics without having to move the data to another platform. For edge analytics applications, this can be a powerful tool to transmit the filtered data to the Cloud or central storage resulting in improved security and reduced cost.

Enterprise Grade Security: Atlas also provides enterprise-grade security measures required for edge applications, including out-of-the-box authorization and authentication methods, TLS encryption during device-to-cloud synchronization, and on-device encryption. These offerings significantly reduce the time it takes to build, deploy, and update data layers for applications at the edge.



Predictive Maintenance

Utilizing sensors and advanced analytics, the health condition of a machine or equipment can be monitored and predicted. This information can then be used to find out the most optimal time for maintenance of the said machine or equipment. Predictive maintenance is a superior form of maintenance when compared with time or usage based maintenance strategies since it promises maintenance cost savings and results in optimized spare part inventory. It also results in better failure forecasting and incipient failure detection, the identification of undesired patterns in data as primary causes of failure. Automotive companies implementing predictive maintenance can analyze reliability data from their global equipment database in order to optimize their overall asset management strategy as well. MongoDB Atlas provides a number of relevant features that help our clients implement predictive maintenance in their shopfloors.

Document Data Model: MongoDB documents are inherently flexible while allowing data governance when required. Since machine health prediction models require not just sensor data but also maintenance history and inventory data, the document model is a perfect fit to model such disparate data sources.

Atlas Search: During the maintenance and support process of a physical product, information such as product information, replacement parts documentation, labor times, damage codes, repair instructions etc. needs to be available at

all times and easily accessible by support staff. Full text search capabilities provided by Atlas can get integrated with the support portal and help staff in retrieving information from Atlas clusters with ease. A major European automotive manufacturer is using Atlas to explore ways of simplifying engine diagnostics. Audio files are recorded from engines which can then be vectorized and searched to retrieve similar cases. Further improving matching, they use a hybrid-search approach. This combines Atlas Search for keyword-based metadata search (i.e. auto model, year of manufacture, plant of manufacture) and audio-based vector similarity retrieval.

Data Lake: As data is ingested, MongoDB Atlas Data Lake automatically optimizes and partitions the data in a columnar format and structure best for analytical queries. This significantly reduces the complexity of transforming data for the data scientist tasked with building the prediction model.

Data Federation: MongoDB Atlas Data Federation provides the capability to federate queries across data stored in various supported storage formats, including MongoDB Atlas Clusters, Data Lake Datasets, AWS S3 Buckets, and HTTP Stores. This feature reduces the complexity of bringing data together for prediction model testing purposes.

Ecosystem Integration: MongoDB Spark Connector opens up access to all Spark libraries for use with MongoDB datasets: Datasets for analysis with SQL (benefiting from automatic schema inference), streaming and machine learning.



Adaptive Quality Assurance

Industry 4.0 technologies enable a step change in quality systems for automotive manufacturers. By leveraging IIoT technologies, it is possible to achieve in-line quality control with traceability and shop floor integration. There are also opportunities to enable adaptive machining with close loop feedback. This means that the machining shop parameters are getting adjusted and optimized in real time as the product quality measurement results are collected and analyzed by the analytical system. Adaptive quality assurance positively impacts Total Cost of Quality (TCQ) for an organization by:

- Increasing labor productivity, advance analytics reduce workload of engineers to conduct basic analysis of data
- Cost savings, early detection and prediction of quality problems and patterns
- Safety, reduce rework for customers and company

MongoDB Atlas provides a number of relevant features that help our clients implement adaptive quality assurance in their shopfloors.

Data API: Companies can use MongoDB Atlas Data API to integrate MongoDB Atlas into any apps and services that support HTTPS requests. Leveraging this feature, the critical measurement results from the Quality Management System can be sent to MongoDB Atlas as needed and then used for analytical purposes using Aggregation Framework or via MongoDB Atlas connectors with third party analytical software.

Online Archive: With MongoDB Atlas Online Archive, manufacturers can push cold data or infrequently accessed data from a MongoDB cluster to a fully-managed cloud object storage. This can lower data management costs for old data, while active data that is more often accessed and queried remains in the primary database. The archived data can be used for audit purposes.

Document Data Model: MongoDB documents are inherently flexible while allowing data governance when required. For adaptive quality assurance use cases, the manufacturer has to not just store machine parameter/config files but also the product dimension measurement reports. The document model is a perfect fit to model such disparate data sources.

MongoDB Atlas Charts: MongoDB Charts is the best way to visualize MongoDB data. Charts is built specifically for the document model, no ETL, no time loss to data manipulation or duplication required to visualize rich JSON data. Using Charts, powerful engaging data experiences can be created for the use case stakeholders in no time.

Atlas Search: Daily production operations generate a vast amount of data especially when machines are IIoT enabled. It is important for the manufacturer to have shop floor applications developed with integrated search capabilities as enabled by Atlas Search that can help operators and supervisors have access to production and product related data (product IDs, Bill of Material, Bill of Process, Defect Lists, Inventory Information) fast.



Virtual Factories

A virtual factory is a computer-based twin of a real-world manufacturing facility. It displays most if not all of the characteristics of a physical factory and allows the manufacturing stakeholders to test out different change or update scenarios before implementing them in the physical factory.

As automotive manufacturers start their digitalization journey, they will start sensorizing their machines and equipment to get data driven insights on their shopfloor manufacturing

processes. They can use this information to optimize their daily production. While a smart factory provides automation and optimization within the physical production environment, a virtual factory complements it by offering simulation capabilities, remote monitoring and control, training and skill development opportunities, and agile operational practices. Together, they enable manufacturers to achieve higher levels of efficiency, productivity, flexibility, and innovation.

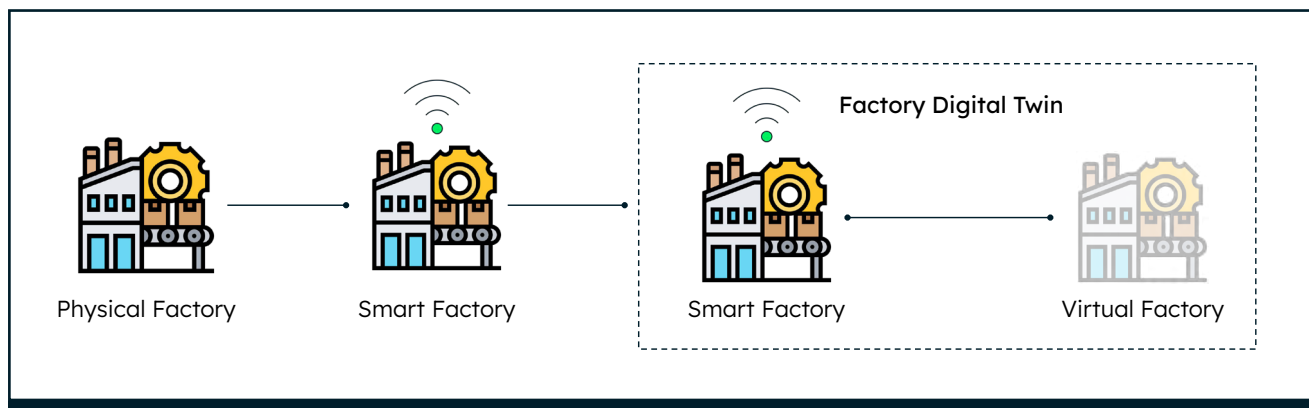


Figure 9. From Smart Factories to Virtual Factories

Virtual factories provide many benefits to automotive manufacturers, some of which include:

1. With the rising inflation across the board, Virtual factories can be used to optimize production processes and identify inefficiencies before they become costly problems. This can lead to increased efficiency, reduced waste, and improved quality.
2. Raw sensor data coming from a physical machine can be contextualized with the actual manufacturing process and then analysis can be run on it to understand the root cause of any quality issues. This can help with implementing quality control while dealing with complex production processes.

3. Using a virtual factory, companies simulate manufacturing processes and test new products or ideas without the need for physical prototypes or real-world production facilities. This can significantly reduce costs associated with research and development, as well as minimize the risk of product failure.

A connected Virtual factory is a relatively new concept. Although discrete event and agent based simulation have existed for many years, companies have started to incorporate connected virtual factories in their digital manufacturing roadmaps. Big players such as [Hyundai](#) and [BMW](#) have announced their intention to build and utilize all the benefits that a virtual factory may provide them.



There are a lot of challenges associated with building a virtual factory. The biggest one is that there are multiple systems and technologies to extract data from, this data needs to be modeled and managed, the manufacturing process is complex and we need a very realistic visualization of the virtual factory. Also as the physical factory grows to be more complex over time with changes to existing layouts, the virtual factory needs to be kept up to date.

To put the virtual factory concept in the perspective of product and production life cycle, when automotive manufacturers wish to create a new product they start with an idea then define the requirements, develop it, produce it and then recycle it. Throughout this time the Product life cycle management system is storing the information along the life cycle of the product. On the other hand, for a production lifecycle, they

have to start with the production requirements, do system engineering and integration and then move the product into production. During this time, the physical factory is generating data and exchanging information between the Enterprise Resource Planning (ERP), Manufacturing Execution System (MES) and the IoT platforms. The virtual factory sits somewhere between product and production life cycle, the data from both life cycles have to flow into the virtual factory. It is critical to have a data platform in the middle that can contextualize all this information coming in and then feed that to the virtual factory and vice versa. MongoDB Atlas acts as that data platform in the middle, providing synchronization capabilities between physical and virtual worlds, enabling flexible data modeling and providing access to the data via a unified query interface. As the data grows, MongoDB can scale horizontally ensuring that all of the data is managed effectively.

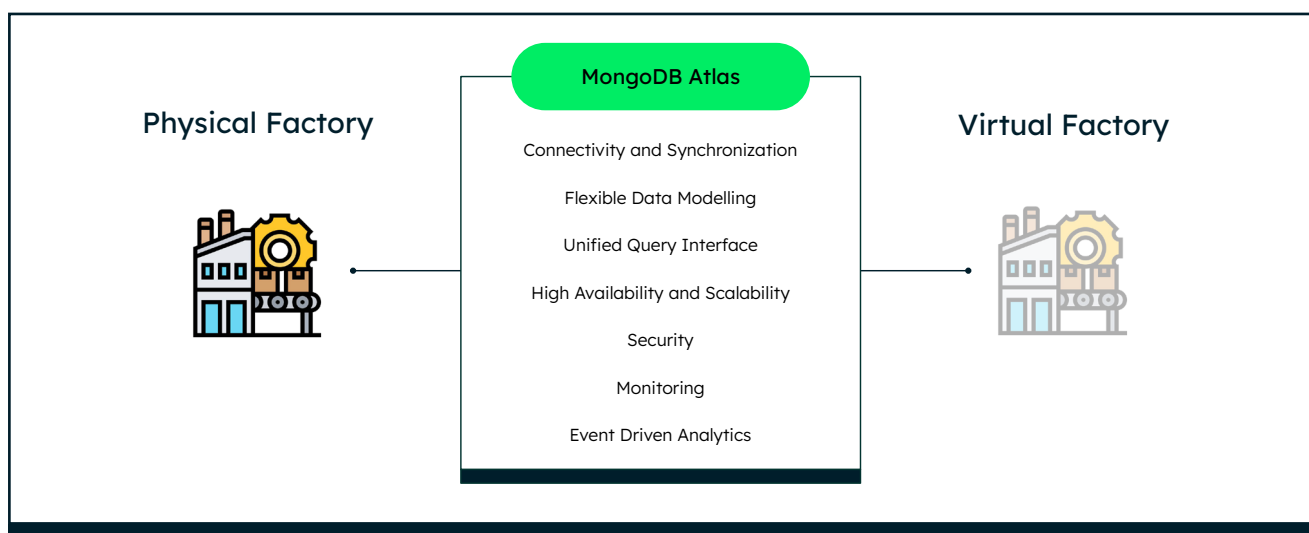


Figure 10a. MongoDB Atlas as the data platform connecting physical and virtual worlds



The virtual factory can be set up as a 3D model of the physical factory in a game engine such as Unity 3D. Using MongoDB Realm as an embedded, object oriented database, virtual factory data and events can be persisted inside the Unity application and then synchronized with MongoDB Atlas through Atlas Device Sync. For the physical factory, the data from machines can be published to an IoT aggregation engine such as AWS IoT Core and from there pushed to Atlas to be stored in either a normal or a time series collection. The data stored in Atlas can be used

to train a machine learning model leveraging our connectors such as Spark or Kafka connector to draw insights on inventory status or production cycle times. MongoDB is strongly positioned to implement virtual factory capabilities for the manufacturing industry, with a strong set of features and functionality that cover the entire lifecycle of manufacturing data. These capabilities allow MongoDB to be in a unique position to fast-track the digital transformation journey of automotive manufacturers.

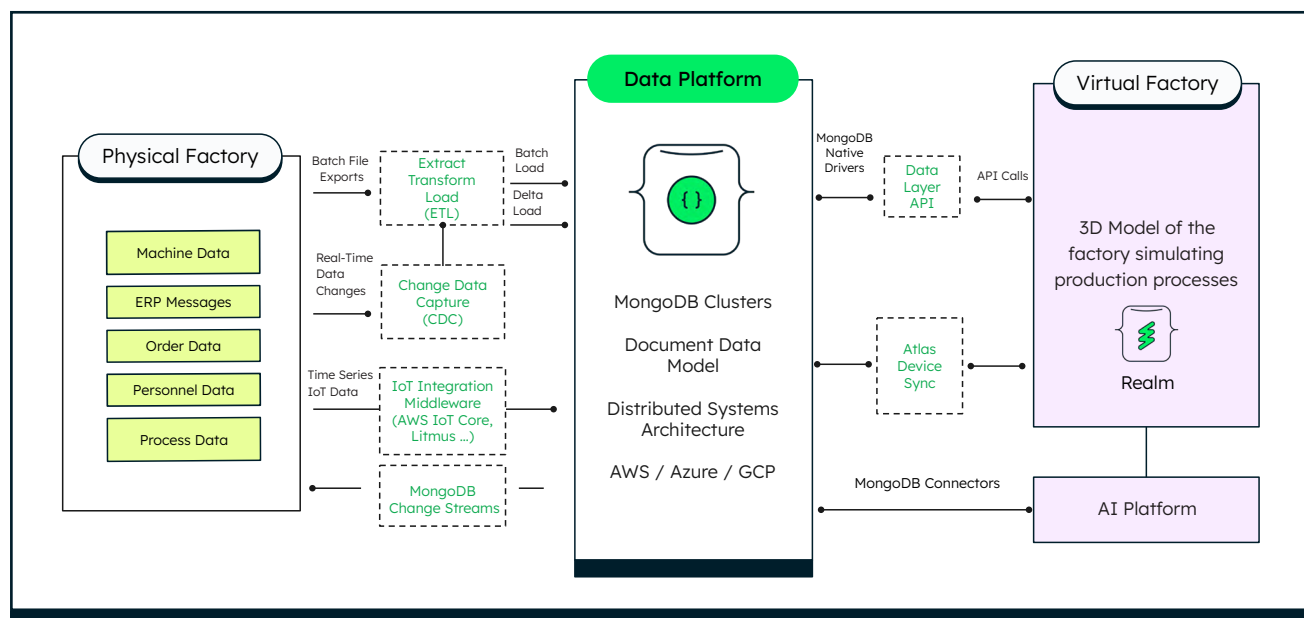


Figure 10b. MongoDB Atlas as the data platform connecting physical and virtual worlds

Smart Commerce

Customer 360

The more automotive manufacturers know about their customers, the better they can design new features in their vehicles and increase the lifetime value of their business. The ability of automotive manufacturers to collect, secure, store and analyze customer data across disparate systems to understand customer behavior is called Customer 360. Having a good customer 360 strategy results in lower Customer Acquisition Costs (CAC) and greater Customer Lifetime Value (CLV). Analytical and AI based models that calculate CAC and CLV need to be fed with up-

to-date contextual and reliable data to identify patterns and derive insights. Creating a holistic 360° view of vehicle customers is a business imperative for every automotive organization. However, creating a unified view or golden record of customers is hard and expensive, especially if you want to capture and make sense of the data from various sources such as ecommerce platforms, Customer Relations Management (CRM), Partner Relationship Management programs, ERP, Customer Service systems, loyalty programs, payment portals, web apps and other related functional and business units.



A Customer Single View solution centralizes and organizes data from all related sources into a standardized format and unified data model. It addresses the problem of disjointed and duplicate data to give a single version of truth of customers. It also includes a consumption component providing a unified view of customers across all related data domains and enables organizations to rapidly evolve their schema and aggregation mechanisms. Finally a customer 360 or single view serves as an ideal foundation for analytics, AI/ML, reporting, and enhancing current operational capabilities.

Traditional Master Data Management systems are typically built on RDBMS having a rigid data model that makes implementation complex, inefficient and slows development. It becomes difficult to store and aggregate data from isolated data sources having diverse structures, different data models and systems holding partial and inconsistent data. They are also often limited by vendor supplied matching algorithms and schema structure is limited to RDBMS platform and vendor defined requirements.

MongoDB Atlas makes achieving Customer 360 easy. Using MongoDB's document model, manufacturers can incorporate any type of data,

no matter what it looks like or where it comes from. JSON documents support all the basic data types (e.g., numbers, strings, binary data, arrays) without a requirement to define or enforce a schema. MongoDB documents can also vary in structure, which means documents from one system don't need to have all the same fields as documents from another. MongoDB's expressive query language, indexing, and aggregation capabilities make it possible to find and filter data, no matter how the business needs to access it.

For data accessibility, multiple features are available. Full text search, graph processing, geospatial, and real time aggregation provide powerful ways to access and analyse customer data. MongoDB distributed, self healing architecture with in-memory processing accommodates always-on and real-time needs. Strong data consistency ensures that the "single source of truth" is unquestionable, even when distributed across many clusters. For data visualization and advanced analytics, plug and play connectivity with SQL based BI tools via Atlas SQL interface and advanced analytics frameworks through Spark connector is available. Finally there is a comprehensive information security model to enforce access controls, encryption, and auditing.

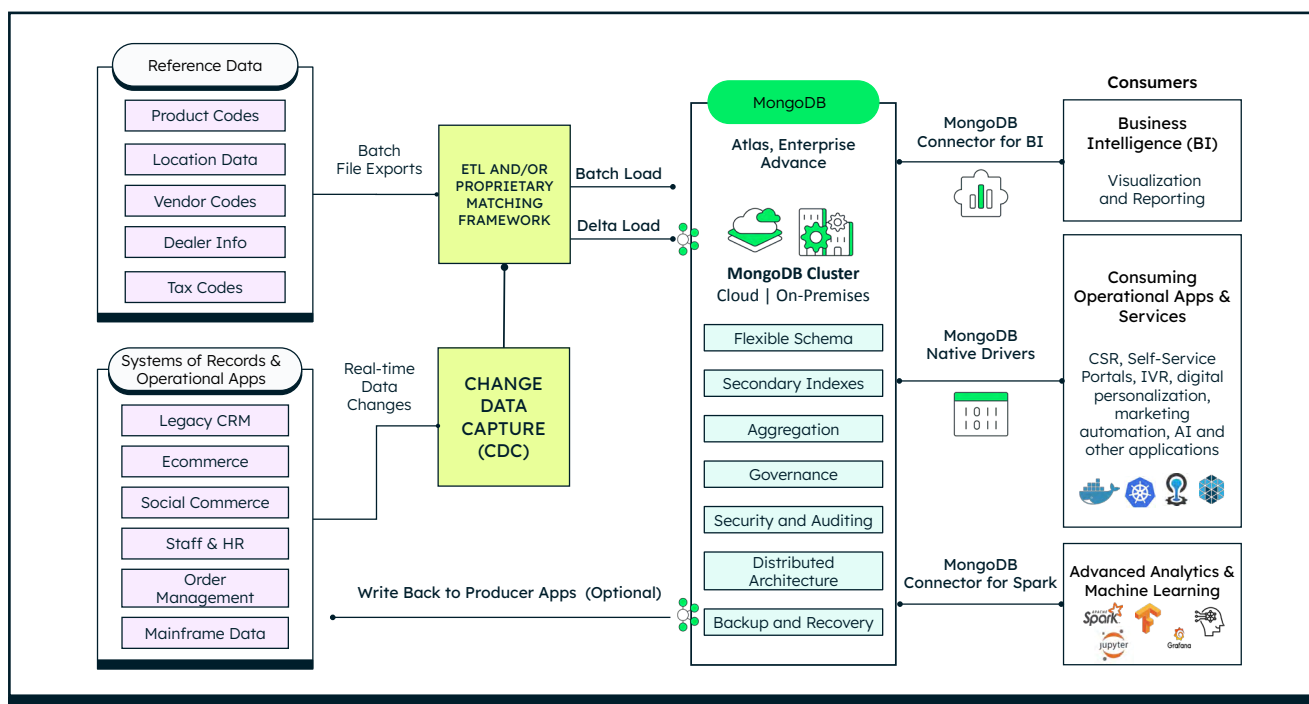


Figure 11. MongoDB Atlas as the Customer 360 data platform



Product Catalogs

The product catalog is at the heart of any automotive retail company- simply put it is the window into the list of all vehicle related products that are for sale and their information. These days that information can be extensive- from name, dimensions, price (and per-store price) to images, videos, and third party reviews. For a seamless digital, in-person or omnichannel experience, the data needs to be consistent and available in real-time for a number of different retail front ends- the ecommerce website, the customer facing application, in-store workforce applications, POS devices and more.

Product catalog data management is a complex problem to solve. Older, more established retailers have traditionally relied on vendor-supplied product information management systems (PIMS) as their backend data store. Product data will be input to the PIMS and then a subset of it will be ETL'd to a product catalog data store to be combined with other necessary information that the viewer will require. For example, the product catalog for ecommerce will need review data and upsell recommendations, whereas the product catalog for the workforce mobile app will need in-store stock location information. The rigidity of the relational database management systems (RDBMS) used often meant that a new product catalog data store was created on top of the PIMS for each use case. This led to a proliferation of data silos that each needed to be updated every time a new product or product information was added, and unreliable consistency when it comes to offering an omnichannel experience as stock and location of items is near impossible to reconcile.

After years of relying on PIMS and multiple monolithic, vendor-provided systems, automotive retailers have learned that product catalogs built on legacy databases are unsuitable for modern ecommerce experiences. There is a chance of loss of revenue and customer trust due to inconsistencies in product data resulting from siloed data or spaghetti architecture that most companies are forced to put into place to bridge the gap between their legacy estate and modern front end applications. There can be loss

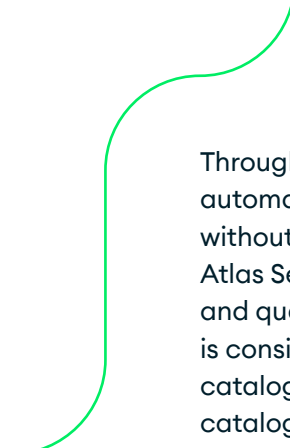
of competitive edge due to delays in new product info going live. Relational data models that require up-front schema definition and downtime to change schema inhibit the ability for new products (with potentially different attributes) to be added or new attributes for existing ones to be discovered. Additionally, the inability to flexibly scale data infrastructure can lead to slow-downs as data size increases.

A modern product catalog provides a single consolidated view of the holistic product catalog with flexibility to model and evolve diverse products over time. It also gives scalability to add products and product data without sacrificing speed and includes a search engine that is easy to enable and provides fast sophisticated results. Finally there are analytics capabilities provided to offer personalized up-sell and cross sell for a variety of products. These are the tenants of a modern operational data layer (ODL): a single consolidated view across all data that has the flexibility, scalability and technical capabilities to serve the operational and analytical needs of the consuming business applications.

Operational use cases are exactly what MongoDB was designed for; the ability to handle high throughput workloads at speed to serve modern application needs. MongoDB's Documents are a natural way to describe data- in product catalog use cases where descriptions of physical objects are stored, the mapping to a document becomes a natural fit. MongoDB's aggregation framework helps create sophisticated processing pipelines for data analytics and transformations. For a product catalog that must serve multiple consuming applications- both customer and business facing - it's imperative that data be easy to work with. Flexibility to model and evolve diverse products over time

In MongoDB, fields can vary from document to document within a single collection. This makes it possible to merge data from source systems storing records on overlapping but non-identical sets of entities. This locality of data ensures the complete document can be accessed in a single database operation that avoids the need internally to pull data from many different tables and rows.





Through horizontal sharding, MongoDB helps to automatically rebalance the data accordingly, without manual intervention. Finally MongoDB Atlas Search allows fine-grained text indexing and querying of data on Atlas clusters. Search is considered table stakes in modern product catalog applications, by implementing a product catalog ODL on MongoDB Atlas, retailers can add functionality like full text search, facets, fuzzy-matching, highlighting, synonyms etc. without considerably increasing their time to market.

Inventory Management

Mobile technologies can be powerful enablers of an improved automotive parts retail experience, both for customers and for employees. Using mobile devices, staffers can check in on an app rather than physically signing into their workplace. Order management and order tracking can be greatly improved with mobile technologies, offering customers increased assurance and predictability.

One of the largest challenges faced by automotive retailers (both parts and vehicles) – and one where mobile technologies are ideally suited to help – is inventory management. Poor inventory management leads directly to lost revenue, because retailers cannot fulfill orders if they don't know what they have. Inventory management is exceptionally complicated, as when an order comes in, the information must be sent to the appropriate warehouses, and the warehouse managers need to know which items should go to which stores. If the various inventory systems aren't communicating well, the store manager may not see that an order for car parts has been returned or canceled. They may think the item is out of stock – even though it could be sitting on a warehouse shelf.

Such difficulties can be the result of fragmented technologies. Some of the inventory data systems in use may be modern, some may be legacy, but without being connected to a central data layer, the information contained within them remains siloed and therefore of limited value.

Mobile technology, connected to a central data layer, can unlock that information and provide important new functionality to employees throughout an organization. It can tell the manager that their parts will arrive in the next hour. And it can tell them that, based on data from the point-of-sale system, they're running out of a specific car accessory, for example, and need to replenish their stock.

The solution is to have a central data layer that is able to talk to all these different systems on a real-time basis and then cascade the information to different applications using mobile devices to feed off of the database. MongoDB Atlas is well suited to be that central data layer with Atlas Device Sync to connect data between mobile devices and the cloud.

The combination of MongoDB Atlas and Atlas Device Sync provides significant inventory management advantages to automotive retail customers, including:

- **Out-of-the-box, robust sync protocol:** When a sale is made, data from the point-of-sale system is synchronized to the cloud and then sent to all the other devices that need the information in real time. No matter how an employee is accessing the data, they can be assured that it is the most up-to-date.
- **Efficient data consumption from the mobile client:** It's important that solution be specifically tailored for mobile devices. A generic cloud-based solution is not necessarily mobile-friendly, which means that it may not be efficient in terms of data consumption, battery-life or storage optimization.
- **Flexible schema:** A flexible data structure allows businesses to ingest data quickly, innovate, and be free from the rigidity of relational systems. Warehouse teams may need one view of the data, while delivery teams need a completely different view. A flexible schema allows for fluid processes and helps each team get the data access it needs.



- **GraphQL and auto-generation of endpoints:** Endpoints talk to the data layer so that an API can fetch data from it. Most developers have to build those endpoints, but with Atlas Device Sync clusters, the endpoints are generated automatically.
- **Triggers and functions:** Certain aspects of MongoDB Atlas and Atlas Device Sync can be

set up to be serverless. Using triggers, data can be moved between Atlas and other systems and send out notifications based on changes in the database.

- **Data access permissions and user authentication:** This functionality is available out of the box with Atlas and Atlas Device Sync, which saves developers valuable time.

Smart Products and Services

Connected Mobility and Digital Customer Experience

There are an estimated 100 million connected vehicles on the road worldwide today, and that number is [expected to reach 400 million](#) by 2025. Connecting vehicles is of utmost importance to OEMs as buying decisions and future growth opportunities are more and more dependent on software capabilities. Many modern vehicles are already in some way connected to the cloud, providing basic capabilities such as sending diagnostics data for predictive maintenance and smooth fleet operation. While many people talk and discuss autonomous vehicles and how long it may take until this becomes a reality, one might overlook all the other, shorter term optimizations and opportunities.

With 5G, bandwidths have dramatically increased and latency is reduced, allowing many more

applications to be built around and within the vehicle. The way this is done today, however, still takes a lot of time due to the amount of non-differentiating “plumbing” efforts, consuming precious time for developers. Speed is crucial to remain competitive, making it imperative to remove the non-differentiating work that’s holding back innovation.

MongoDB helps remove the plumbing work and increases developer efficiency by providing an end-to-end data platform for data processing inside the vehicle, the cloud, and mobile seamlessly integrated and purely consumption based on the hyperscaler of your choice.

MongoDB Atlas, the developer data platform on AWS, Azure, and GCP has three core capabilities, helping you to build faster across different environments to power wide varieties of fleet management applications or innovative digital customer experiences.

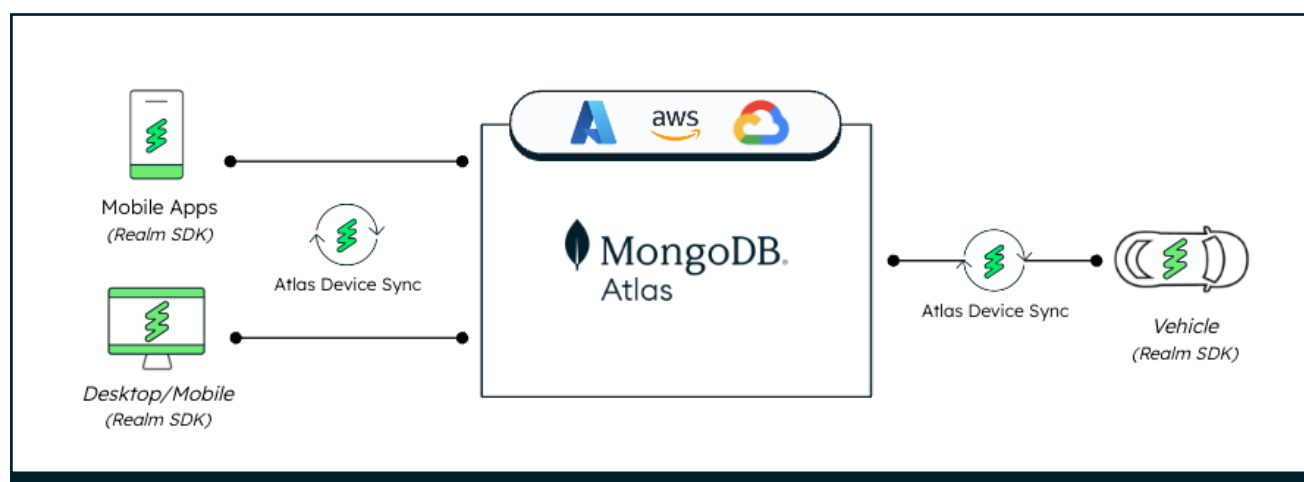


Figure 12. MongoDB Atlas as the Connected Vehicle platform



MongoDB Atlas: Cloud Data Backend

Millions of vehicles and large bandwidths ultimately lead to large volumes of data and high expectations towards reliability and high availability for any kind of digital service tied into it. At the core of MongoDB Atlas is the hugely popular MongoDB document model database technology.

- **Document Model:** A document is structured like the broadly adopted and popular JSON structure. The document model is the superset of the other data models and allows you to seamlessly store and process key-value pairs, geospatial, tabular and graph data together with time series in a single data platform and through a single unified API and query language.
- **High Availability & Scalability:** MongoDB is built with a distributed architecture from the ground up. While you can run a single node, the typical setup is a three node replica set providing automatic and transparent fail overs from one node to another and thus allow you to run maintenance and rolling upgrades while the applications are fully operational without interruption. You can add up to 50 nodes to a replica set and distribute those nodes across

Use Case

SHARE NOW, a car-sharing service, is disrupting the industry with a new way of car rental in Europe. SHARE NOW uses data and the Internet of Things (IoT) to carry out centralized management of its vehicles. But the volume of data was becoming a developer burden. Using MongoDB Atlas, SHARE NOW's developers no longer need to worry about scaling, provisioning, writing, and configuring its 200 microservice document databases. Instead, the team can focus on building new application features and improving services much faster for its customers – wherever they may be on their travels.

[Read the full story.](#)

the globe for low latency access and read scalability. On top of that, MongoDB has built-in sharding to scale writes at global scale while providing full control over data locality to accommodate regulatory requirements.

Realm SDKs: Embedded Lightweight Database

Vehicles and smart phones have a few things in common, including limited battery capacity and compute resources. A few more MB memory in a car multiplied by millions of cars quickly costs a couple of million dollars, not to mention that more compute drains the battery faster. Therefore, it's crucial that any application deployed into such a constrained environment is as resource efficient as possible.

Opposite to MongoDB, Realm is an object oriented, embedded database with no separately running process. As a library it becomes directly part of your application code and can be easily deployed through existing procedures and tooling. It's object orientation is idiomatic to software developers, helps reduce complexity by removing ORM libraries and is available as C++, Flutter, Swift, Kotlin, JavaScript, and C# SDK.

Atlas Device Sync: Bi-directional Data Synchronization

Powerful data processing capabilities on the cloud backend and lightweight data management for mobile and vehicle applications combine to power greater developer efficiency. But there is still a lot of plumbing required to connect these applications across different environments. MongoDB Atlas solves this problem by providing bi-directional data synchronization efficiently across unreliable networks.

Mobile and vehicle applications store the data inside of Realm files and allow usage of the application even while being offline. Any changes to the data are stored locally first, and as soon as connectivity is established, it is synchronized with the cloud backend and other connected applications. Potential conflicts are deterministically resolved for maximum software developer efficiency around building differentiation.



Conclusion

MongoDB has emerged as a pivotal technology driving innovation within the auto industry. By providing a flexible and scalable platform for handling vast amounts of data, MongoDB empowers automotive companies to unlock valuable insights, enhance customer experiences, and optimize operations.

With its document-oriented data model, real-time analytics capabilities, and seamless integration with other cutting-edge technologies, MongoDB has become the go-to choice for forward-thinking automakers and suppliers. As

the industry continues to evolve and embrace digital transformation, MongoDB's role in fueling innovation is set to expand, enabling automotive companies to stay ahead in a rapidly changing landscape and shape the future of mobility.

About the Author



Humza Akhtar, Ph.D. is a Principal on the Industry Solutions Team at MongoDB focussed on manufacturing and IoT use cases. Prior to joining MongoDB, he was working at Ernst & Young Canada as a Senior Manager, in digital operations consultancy practice. Humza attained his Ph.D. at Nanyang Technological University, Singapore, and worked with the Singapore manufacturing industry for a number of years on Industry 4.0 research and implementation. He has spent most of his career enabling smart and connected factories for many manufacturing clients. In 2020-2021, he established a multi-year strategic roadmap for Singapore's smart supply chain initiatives. His book on Industry 4.0 is the first book of its kind detailing the real-world implementation of concepts related to digital manufacturing.

[Learn more](#) about MongoDB for Automotive Manufacturing.

