




MongoDB Atlas for
Industries

Manufacturing and Motion

The integration of AI within the manufacturing and automotive industry has transformed the conventional value chain, presenting a spectrum of opportunities.



AI investments in the manufacturing industry



AI in manufacturing is a game-changer. It has the potential to transform performance across the breadth and depth of manufacturing operations. Companies are in a race to embrace AI, as these technologies are critical enablers of the Fifth Industrial Revolution (also known as Industry 5.0). Artificial intelligence in manufacturing is bringing factories into the future and will ultimately empower the manufacturing market to continue to be the backbone of the global economy.

Industry-wide, manufacturers are facing a range of challenges that make it difficult to speed production while still providing high-value and high-quality products to their customers. All the while, companies need to implement a digital infrastructure that positions them to fully embrace the skills and knowledge of their best assets — people.

The manufacturing industry today relies on automation just as much as people. But the factory of the future, which is a marriage of physical and digital capabilities, requires more: real-time data, connectivity and AI technology at the forefront. In fact, **more than 80%** of C-suite executives believe they must leverage AI to achieve their growth objectives. Customer requirements for delivering on-time and on-budget product are of the utmost importance, and efficiency is a goal in everything.

AI's ability to drive impact in this regard is real. According to a study of Capgemini, three use cases stand out in terms of their suitability for kickstarting a manufacturer's AI journey: demand planning, intelligent maintenance and product quality control. These use cases have

an optimal combination of several characteristics, that make them an ideal place to start:

- Clear business value/benefits
- Relative ease of implementation
- Availability of data

Smart manufacturing use cases are revolutionizing many organizations, and a key driver of this is the incorporation of artificial intelligence into manufacturing processes. Many firms have embarked on significant digital transformation journeys in the past two years with the goal of improving efficiency and resilience. However, a concerning gap exists between tech adoption and return on investment. While 89% of organizations have begun digital and AI transformations, only 31% have seen the expected revenue lift, and only 25% have realized the expected cost savings (McKinsey). In some cases, situations have even worsened.

This highlights the importance of not just implementing new technologies, but implementing them strategically. In other words, simply deploying AI isn't a guaranteed path to success. Manufacturers need to carefully consider how AI can address their specific challenges, and then integrate it into existing processes effectively.

This chapter unpacks how major players in the manufacturing industry are leveraging AI to improve operations, deliver better outcomes for customers, and realize innovation. It delves into three high impact value drivers and AI use cases: Predictive Maintenance, Inventory Management and Knowledge Management.

The path to success



Successful organizations exhibit common traits across five key areas:

- **Identifying high-impact value drivers and AI use cases:** Efforts should be concentrated on domains where artificial intelligence yields maximal utility rather than employing it arbitrarily.
- **Aligning AI strategy with data strategy:** Organizations must establish a strong data foundation with a data strategy that directly supports their AI goals.
- **Continuous data enrichment and accessibility:** High-quality data, readily available and usable across the organization, is essential for the success of AI initiatives.
- **Empowering talent and fostering development:** By equipping their workforce with training and resources, organizations can empower them to leverage AI effectively.
- **Enabling scalable AI adoption:** Building a strong and scalable infrastructure is key to unlocking the full potential of AI by enabling its smooth and ongoing integration across the organization.

Inventory Management and Optimization

Current State and Challenges

Modern manufacturing supply chains are complex systems, interconnected across the globe. Efficient supply chains are able to control operational costs and ensure on-time delivery to their customers. Inventory optimization and management is a key component in achieving these goals. While maintaining higher inventory levels allows for suppliers to deal with unexpected fluctuations in demand, they come with higher inventory holding costs that may be passed on to customers. Thus, every player in the supply chain is motivated to strike a

balance between inventory levels to maximize profitability and competitive advantage in the market. Effective inventory management mitigates the risk of 'bullwhip effect', where sudden demands can disrupt the supply chain costs and performance.

Key components of supply chain management include procurement and sourcing, manufacturing and production, distribution, logistics and retail. Technological advancements including IoT and AI (including Gen AI) are being integrated into SCM to improve

transparency, efficiency and adaptability of the supply chain, allowing for real time monitoring, predictive analytics, and enhanced decision-making capabilities.

technology-driven ecosystem that requires collaboration throughout the supply chain

between OEMs, tier1-n suppliers and customers, always aiming for reduced costs, quicker production and response times and heightened customer satisfaction, all of which will result in stronger market position.

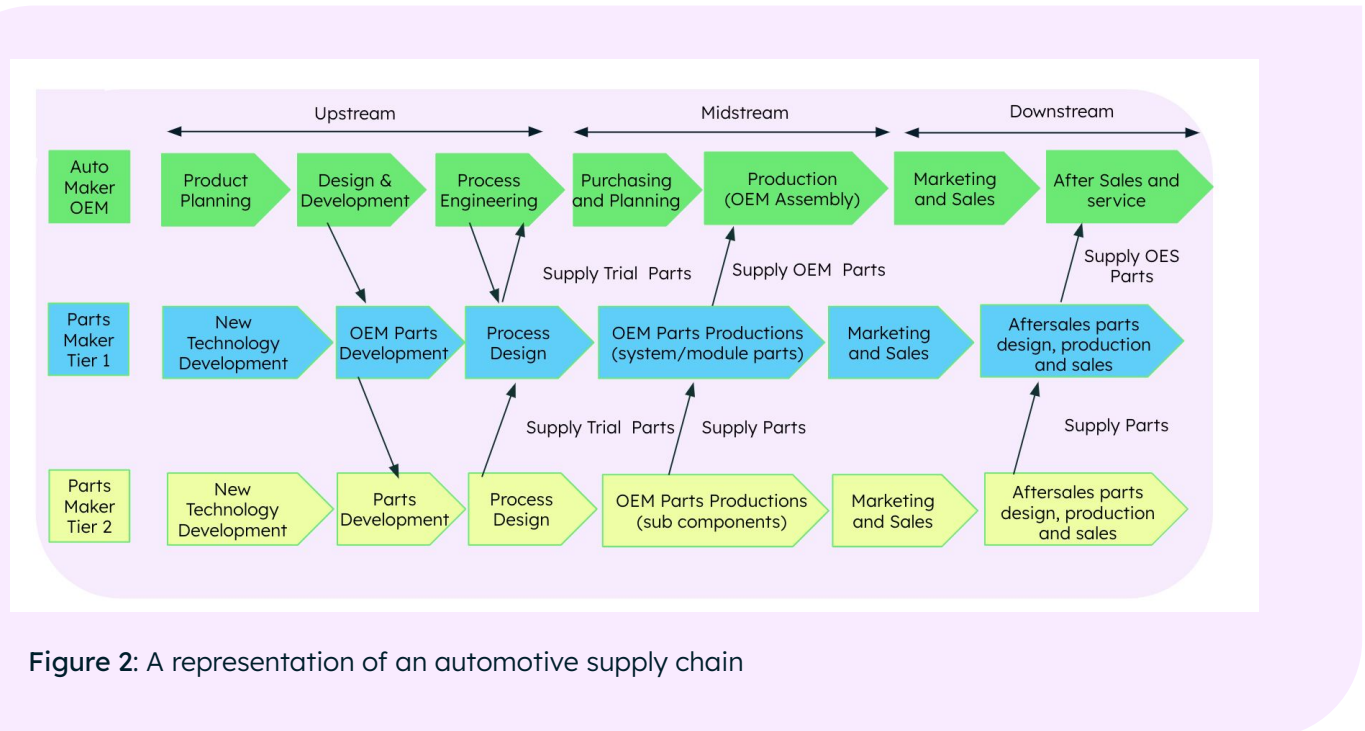


Figure 2: A representation of an automotive supply chain

Inventory management is essential for achieving the goals of efficient supply chains, controlling costs, and delivering to customers with minimal delays. Inventory management is primarily concerned with the planning and controlling an industry's inventory. It includes business processes such as estimating material requirements at various points in the supply chain, determining necessary material's amount, ordering frequency, and safety stock levels. It also includes inventory visibility, inventory forecasting, lead time management, inventory shipping costs, inventory valuation, forecasting future inventory prices, available physical space, quality management, returns and defective goods, and demand forecasting. It plays a very important role in reducing overall costs and rapid response objectives. Effective inventory management requires the

right inventory in the right place at the right time to minimize system costs and meet customer needs.

Usually, companies do supply chain planning at several levels, namely strategic, tactical, and operational. Each level differs in its objectives, planning horizon, and level of detail. Strategic and tactical planning are crucial to successful supply chain management. The so-called Pareto's law could be applicable here as 20% of efforts in strategic and tactical planning brings 80% of the total effect. At the strategic level, the leadership team makes high-level decisions that affect the entire organization. Scenario planning is done at this level. The analysts go through scores of internal and external data including global news, political developments, think tank studies and scientific

literature to pinpoint the most strategic concerns and trends that the organization needs to focus on. The team can then use these outputs to develop a set of draft scenarios for consideration.

This tedious process comes with its own set of challenges. Predicting long-term demand, market trends and economic conditions is challenging because of the long-term horizon. This long planning horizon increases the uncertainty in predicting demand, as market conditions, consumer preferences, and technological advancements can change significantly over time.

At the tactical and operational level, for 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.

However, efficient inventory management for manufacturers presents complex data challenges too, primarily in forecasting demand accurately and optimizing stock levels. One issue routinely faced is the variability and unpredictability of customer demand patterns, making it difficult to precisely anticipate inventory needs. Managing diverse data streams from sales records, production schedules, supplier information, and market trends poses a considerable data

integration challenge. The spread of data across multiple systems and locations (on-prem systems, cloud regions etc.) can lead to data silos and hinder visibility into overall inventory levels and movements. Finally, when there is sparse historical inventory data available, then traditional ML models may suffer in accuracy.

Generative AI and IoT technologies hold potential to address some of these challenges. Generative AI in particular can assist in scenario planning by generating various potential outcomes based on a wide range of data, allowing the organization to prepare better for an uncertain future.

How AI and MongoDB Help

We will start with scenario planning to generate hypothetical situations which could affect inventory requirements, supply chain performance and overall operations. Effective scenario planning helps companies plan for optimal inventory levels.

After scenario planning, we will look into inventory classification use cases where AI can be used to categorize inventory based on factors such as demand variability, lead times and criticality.

Scenario planning

A scenario planning process has two stages

- Scenario generation
- Strategy generation

In scenario generation, a generative AI application can look at a vast amount of data including internal and external business data, competition data, political news and events and social media news, find the correlation between each piece of unstructured information and then rank these areas of concern in terms of their estimated

significance. It is important to utilize general knowledge that a Large Language Model processes as well as internal company data in a Retrieval Augmented Generation (RAG) model to avoid hallucinations.

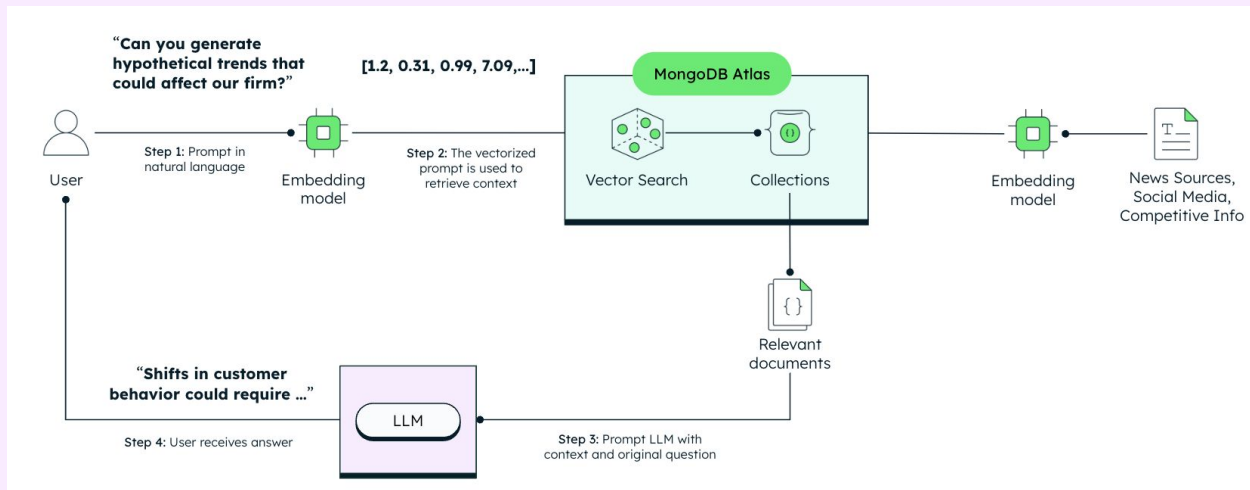


Figure 3: A Scenario Generation RAG App with MongoDB

The figure above shows a scenario generation application data flow. This retrieval-augmented generation (RAG) system consists of four parts. First, an AI data store aggregates and operationalizes structured and unstructured data. In our case, the majority of the data is in unstructured format such as news sources, social media and competitive information typically stored as unstructured PDF files. The PDF files are vectorized and stored in MongoDB Atlas. Atlas Vector Search is then utilized to perform semantic searches and to find meaningful context from the PDF embeddings.

Atlas Vector Search can be triggered using an AI application, connecting to MongoDB Atlas to retrieve the right context, which is then fed into the large language model to answer questions like "Can you generate hypothetical trends that can affect our firm?" The response might suggest customer behavior shifts or economic

factors due to certain reasons, including but not limited to the political landscape and global supply chain disruptions etc.

MongoDB Atlas streamlines RAG implementation as it handles everything under the hood. App data, metadata, context data, and vectors are all stored in the same place. As the app evolves, the document model is inherently flexible and ideal for storing structured and unstructured data. You can add data to the same collections inside the database as needed, vectorize it, and store the vectors alongside it.

Once the data is stored, vector search capabilities are provided right out of the box, and search operations can be optimized using dedicated search nodes.

With MongoDB Atlas, it's just one query in one technology, one language, and one infrastructure to manage and scale, with no data duplication, ultimately leading to a lower total cost of ownership and a unified developer experience.

The same application can then categorize the trends in terms of their probability of occurrence and impact. Next, a strategy can be created automatically, which provides guidance on what steps to take in response to the trends and scenarios generated. The strategy can also contain information about the financial implications and risks associated with the response, for example any significant initial investment in hiring more AI scientists to develop RAG applications for inventory management.

For scenario planning, generative AI can be used to generate and evaluate strategies. However, it is important to provide enough context to the LLM so that it does not hallucinate. MongoDB Atlas Vector Search is key to creating a RAG application. Additionally, these AI-generated strategies should be viewed as initial concepts for further exploration rather than as final solutions to be adopted without additional analysis. Generative AI should complement human efforts by supporting the identification, evaluation and timely execution of appropriate strategies.

Inventory Optimization

One of the most significant applications of AI in inventory management is in demand forecasting. AI algorithms can be used to analyze complex datasets to predict future demand of products or parts. Improvement in demand forecasting accuracy is crucial for maintaining optimal inventory levels. AI-based time series forecasting can assist in adapting to rapid changes in customer demand.

Once the demand is known, AI can play a pivotal role in stock optimization. By analyzing historical sales data, market trends, manufacturers can determine the most efficient stock levels. AI systems can also place orders automatically based on predicted demand and targeted stock levels. This automation not just saves time but also reduces human error. Finally, AI utility can be extended to supplier selection and relationship management. By analyzing supplier performance data, AI based software can assist in choosing the right suppliers who will meet the company's quality, delivery and cost requirements.

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 can be persisted in a MongoDB collection. 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. This repository becomes the source of data for the inventory management AI applications. This approach removes data silos and improves visibility into overall inventory levels and movements.

However, the challenge of poor or sparse data at the source systems may still remain. To solve this, manufacturers can take advantage of generative AI and Atlas Vector Search to implement a Retrieval Augmented Generation (RAG) architecture to generate synthetic data whenever needed. They can take multimodal content such as product descriptions and

specifications, customer feedback and reviews and inventory notes, vectorize them and store vector embeddings alongside the operational data in MongoDB Atlas. This allows them to supercharge their inventory optimization using

RAG. They can easily categorize products based on their seasonal attributes, cluster products with similar seasonal demand patterns and provide context to the foundation model to improve the accuracy of synthetic inventory data generation.

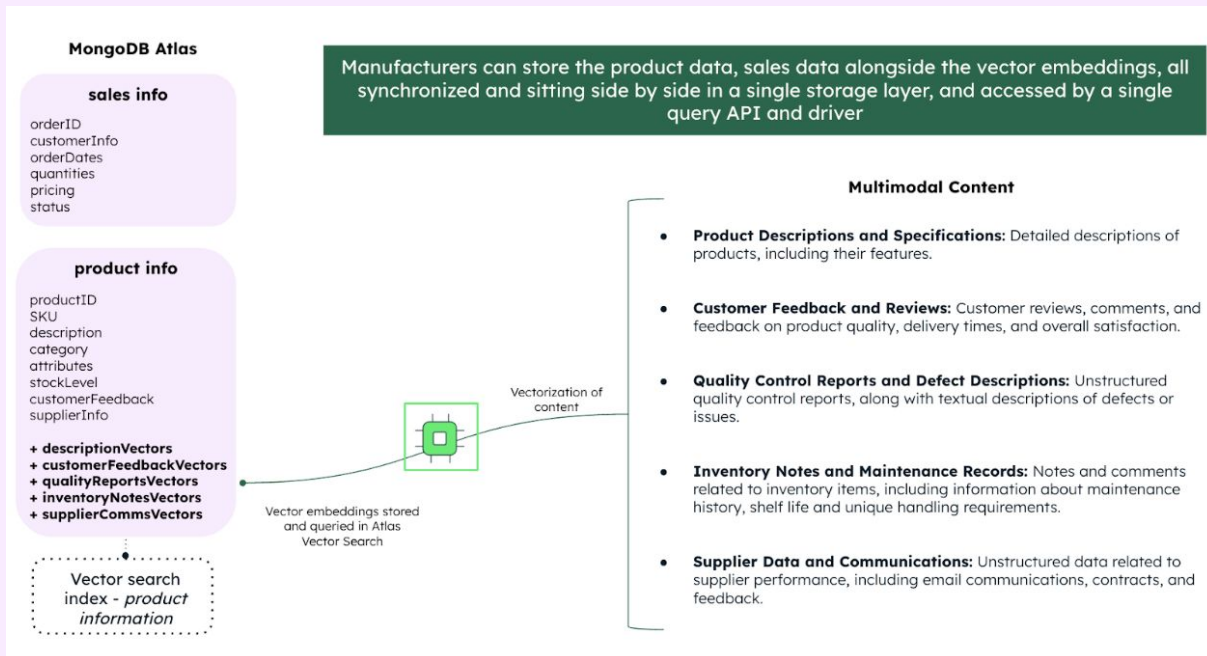


Figure 4: Enhancing Inventory Data with Vector Embeddings

The figure below shows a reference architecture of generative AI+AI enabled demand forecasting with MongoDB Atlas. The accurate demand forecasting results will help in stocking up on right inventory levels. For new products, the historical sales data is not available. Generative AI models can create realistic and diverse synthetic data by learning patterns from existing datasets of similar products. This synthetic data can mimic the sales trends and seasonality that new products might experience. Atlas Vector search can find similar products attributes and feed that context into the generative AI model. By finding semantics in similar products, Atlas vector search can help refine the synthetic data generation, ensuring that it closely reflects potential market conditions and

customer behaviors. This approach not only fills the gap of missing historical data for new products but also provides a robust foundation for demand forecasting enabling manufacturers to optimize their inventory levels.

Solution demo

Discover* how to building an event-driven inventory management system.

*mongodb.com/solutions/solutions-library/event-driven-inventory-management

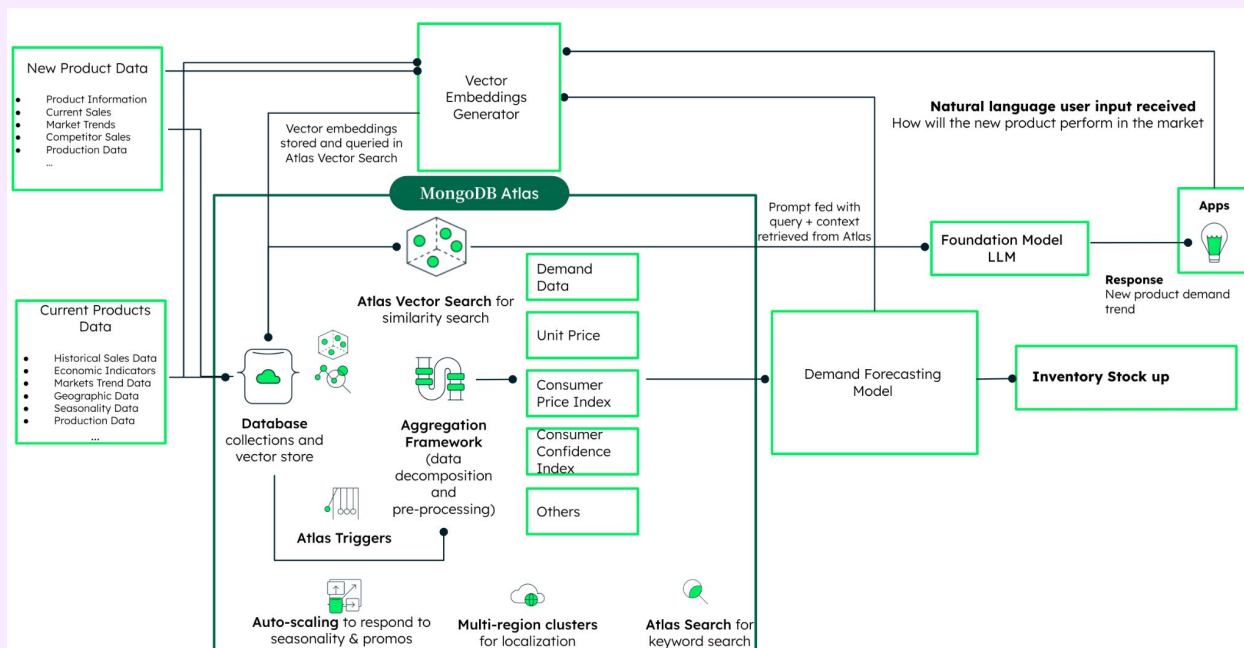


Figure 5: Gen AI enabled demand forecasting with MongoDB Atlas



Ceto is revolutionizing maritime operations with MongoDB time series

Ceto is on a mission to bring the maritime industry into the digital age—and to transform maritime operations into a model of efficiency and sustainability.

To make its mission a reality, Ceto partnered with MongoDB, leveraging its robust data handling capabilities to integrate AI with real-time data collected from thousands of sensors across its customers' fleets. This allows Ceto to predict and preempt potential failures, streamline operations, and manage risks proactively. This shift not only enhances safety and reliability but also propels maritime logistics into a new era of technological advancement, making Ceto a transformative force in global commerce.

MongoDB's architecture provided Ceto with several key features that are crucial for their operations. Scalability was essential for managing the increasing data volumes generated by their expanding fleet. **Time Series Collections** offered advanced data compression capabilities, crucial for managing the large volumes of data generated daily.

“MongoDB's Time Series collections have revolutionized how we manage and utilize data from our fleet. The ability to process and analyze data in real-time has significantly enhanced our predictive maintenance capabilities.”

Learn more*

Ben Harrison
CTO, Ceto

*mongodb.com/solutions/customer-case-studies/ceto

Predictive Maintenance

A well-defined maintenance strategy can be a game-changer for any organization, driving significant revenue and cost savings. Here's how it works:

First, identify the equipment that is most crucial for your operations. Downtime for this equipment can lead to bottlenecks, halting production.

Second, equip these critical assets with sensors to enable condition monitoring. This allows you to monitor the health of the equipment in real time, identifying potential issues before they escalate into catastrophic failures.

Third, based on the prediction, the system can generate work orders, schedule maintenance activities, and even provide guidance to maintenance personnel. This ensures that maintenance is performed only when necessary, optimizing resource allocation.

This series of activities delivers tangible benefits. Costs are reduced through saved labor hours and extended machine lifespan. Additionally, revenue increases as your machines operate at optimal performance levels.

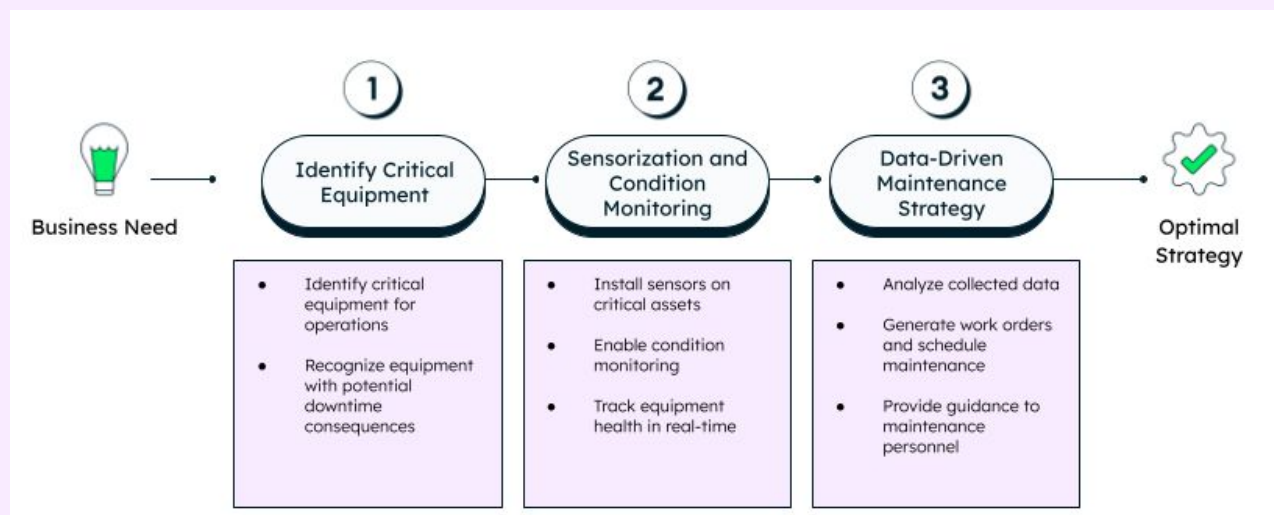


Figure 6: Steps required for an optimal maintenance strategy

Optimal Maintenance Strategy

An optimal maintenance strategy isn't a one-size-fits-all solution. It's about choosing the right blend of approaches based on your specific equipment and operational needs.

In today's processes, we see a spectrum of maintenance approaches. These methods range from highly complex and expensive at one end to simpler, more affordable options on the other:

- **Reactive maintenance:** This is the most basic approach in which maintenance is performed after a machine fails. While simple, it can lead to unexpected downtime and higher repair costs.
- **Preventive maintenance:** This is a proactive approach that involves scheduling maintenance tasks based on predetermined time intervals or usage metrics. This helps prevent breakdowns but can be inefficient as machine conditions can vary. Thresholds for these tasks may need to be adjusted due to factors like aging equipment, changes in processes, or different materials being used.
- **Condition-based maintenance (CBM):** This approach continuously monitors the health of the machine through sensors and data analysis. Maintenance is then triggered based on the actual condition of the equipment rather than a set schedule. This is more efficient than preventive maintenance as it avoids unnecessary maintenance. Threshold-based alerting systems are often used with CBM.
- **Predictive maintenance:** This is the most advanced approach, using data analytics to predict potential equipment failures before they occur,

which allows for proactive maintenance and minimizes downtime. Predictive maintenance requires significant upfront investment in sensors and data analysis tools.

Predictive maintenance uses data analysis to identify problems in machines before they fail. This allows organizations to schedule maintenance at the optimal time, maximizing machine reliability and efficiency.

Here's how predictive maintenance can benefit manufacturing operations, [according to Deloitte](#):

3-5% Reduction in new equipment costs

5-20% increase in labor productivity

15-20% reduction in facility downtime

10-30% reduction in inventory levels

5-20% reduction in carrying costs

Predictive maintenance is constantly evolving. We've moved beyond basic threshold-based monitoring to advanced techniques like machine learning (ML) models. These models can not only predict failures but also diagnose the root cause, allowing for targeted repairs.

The latest trend in predictive maintenance is automated strategy creation. This involves using AI to not only predict equipment breakdowns but also to generate repair plans, ensuring the right fixes are made at the right time.

Automated strategy creation requires substantial investment in R&D, along with deep industry knowledge, access to relevant data, and practical operational experience. The question is, can generative AI help?

Current State and Challenges

The answer is yes, generative AI can help. But there are challenges at each stage of implementation that organizations must consider. Each stage involves a key question and associated challenges, highlighting the steps and issues faced in predictive maintenance and machinery upkeep.

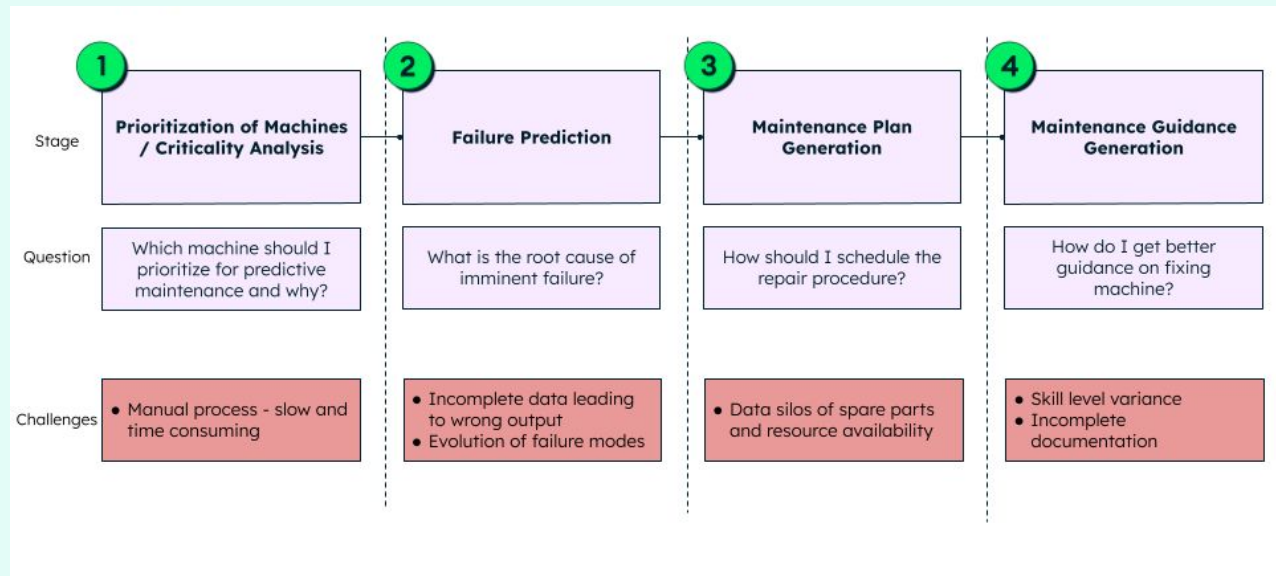


Figure 7: Different challenges seen at each stage of the predictive maintenance deployment

Now, let's envision a factory containing both automated and manual machines. Let's call it Gear Transmission Systems Ltd, whose primary output is gearboxes.

Within this factory, we have an array of equipment: cutting machines, milling machines, measurement devices, and more. As its general manager, you are tasked with managing the budget allocated for maintenance processes and improving strategies. One pressing question you must address is: which machines should take priority for the predictive maintenance projects, and why? This involves consulting with the

maintenance managers and leaders and conducting quantitative analyses, a rather manual process.

Once you've identified the machines, the next step is to install sensors and train the machine learning model. However, two major challenges arise. First, you lack sufficient "run to failure" data to effectively train the model. Secondly, machine health deteriorates over time, leading to evolving failure modes with the age of the machine.

Assuming you manage to overcome these hurdles, the next phase involves maintenance

scheduling and execution. You're faced with a myriad of data silos, including inventory data and resource availability data, which need to be integrated to formulate a comprehensive repair plan. Furthermore, it's essential to ensure that operators are adept at addressing minor machine issues to reduce reliance on external experts. While complex issues may still require OEM or SI support, internal troubleshooting capabilities are invaluable. Therefore, developing easy-to-follow documentation tailored to the skill levels of our staff is important.

As the Figure below shows, different data is required for solving above listed challenges:

- **Prioritization of machines/criticality analysis:** At this stage, we require both structured and unstructured data. We need previous machine failure data as well as expert analysis/opinion on which machines to prioritize for predictive maintenance and why.
- **Failure prediction:** This stage involves structured data such as sensor data and maintenance logs to identify the root cause of imminent failure.
- **Maintenance plan and guidance generation:** In both of these stages, we deal with both structured and unstructured data. The objective is to combine this data to generate an optimal repair plan and operator guidance.

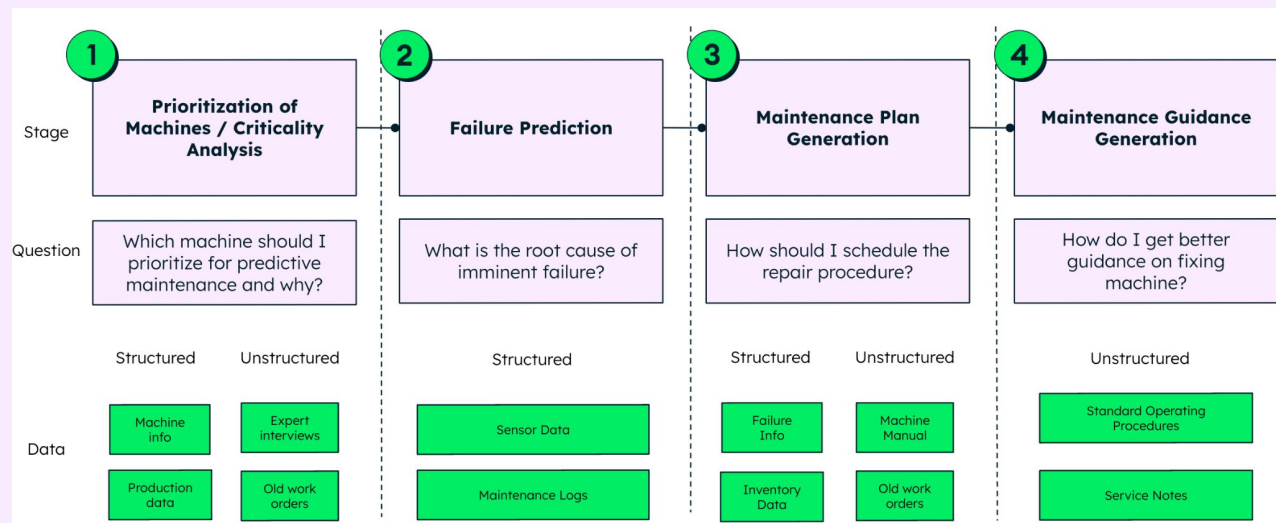


Figure 8: Different data requirements at each stage

How AI and MongoDB Help

MongoDB Atlas is the only multi-cloud developer data platform designed to accelerate and simplify how developers work with data. Using MongoDB Atlas, developers can power end-to-end value chain optimization with AI/ML, advanced analytics, and real-time data processing for innovative mobile, edge, and IoT applications.

Stage 1: Machine prioritization

Current machine prioritization for predictive maintenance relies heavily on manual analysis. Factory personnel gather historical and current machine data on utilization losses due to breakdowns. This data is then reviewed alongside the experience of maintenance managers and leaders. Based on this combined analysis, a roadmap for the predictive maintenance project is recommended, highlighting which machines should be prioritized.

However, this approach has limitations. A reliance on manual analysis can be time-consuming and may not always capture the full picture of the maintenance project due to the limited use of quantitative data sources. Additionally, inconsistencies in interpretation can lead to an overdependence on institutional knowledge, which in turn can result in false analyses that impact the project's return on investment (ROI).

But with the arrival of generative AI, things have changed. A generative AI-based machine prioritization tool can be created to reduce the time manufacturing experts spend on manual analysis, and to decrease the risk of poor investments. To leverage AI, experts need a data store capable of storing and operationalizing both structured and unstructured data. Having such a data store will allow them to perform semantic searches

and to provide the right context to the large language model, ensuring it generates responses based on factory data without hallucinating. Such a system can result in positive business outcomes. Here's how the system can look with MongoDB Atlas as the AI data store:

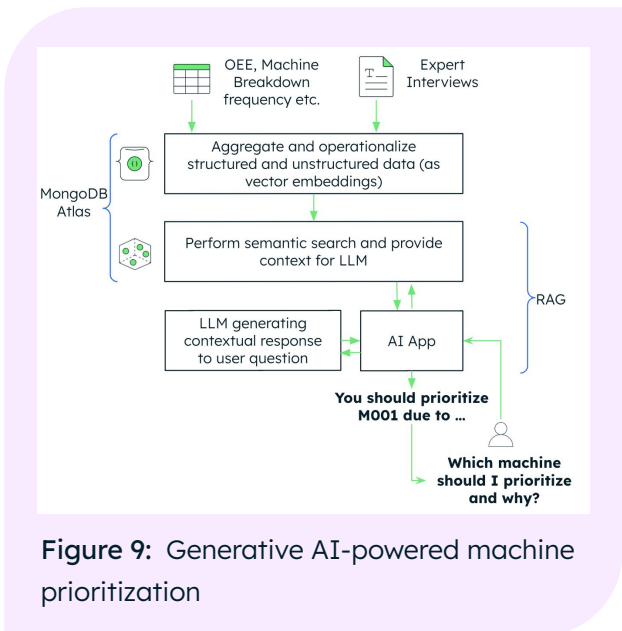


Figure 9: Generative AI-powered machine prioritization

This retrieval-augmented generation (RAG) system consists of four parts. First, an AI data store aggregates and operationalizes structured and unstructured data. In the Figure above, machine breakdown history and operational parameters are represented as structured data, while expert interviews are stored as unstructured PDF files. The PDF files are vectorized and stored in MongoDB Atlas. Atlas Vector Search is then utilized to perform semantic searches and to find meaningful context from the PDF embeddings.

Atlas Vector Search can be triggered using an AI application, connecting to MongoDB Atlas to retrieve the right context, which is then fed into the large language model to answer questions like "Which machine should I prioritize and why?" The response might suggest prioritizing Machine M001 or M002 due to certain reasons, including but not limited to the criticality of the machine, high maintenance cost, etc.

Stage 2: Failure prediction

Now that we've discussed prioritizing equipment, let's move on to failure prediction. MongoDB Atlas provides all the necessary building blocks or tools to implement failure prediction. By providing a unified view of

operational data, real-time processing capabilities, integrated monitoring and alerting, and seamless compatibility with machine learning tools, MongoDB Atlas enables organizations to optimize machine performance and minimize downtime.

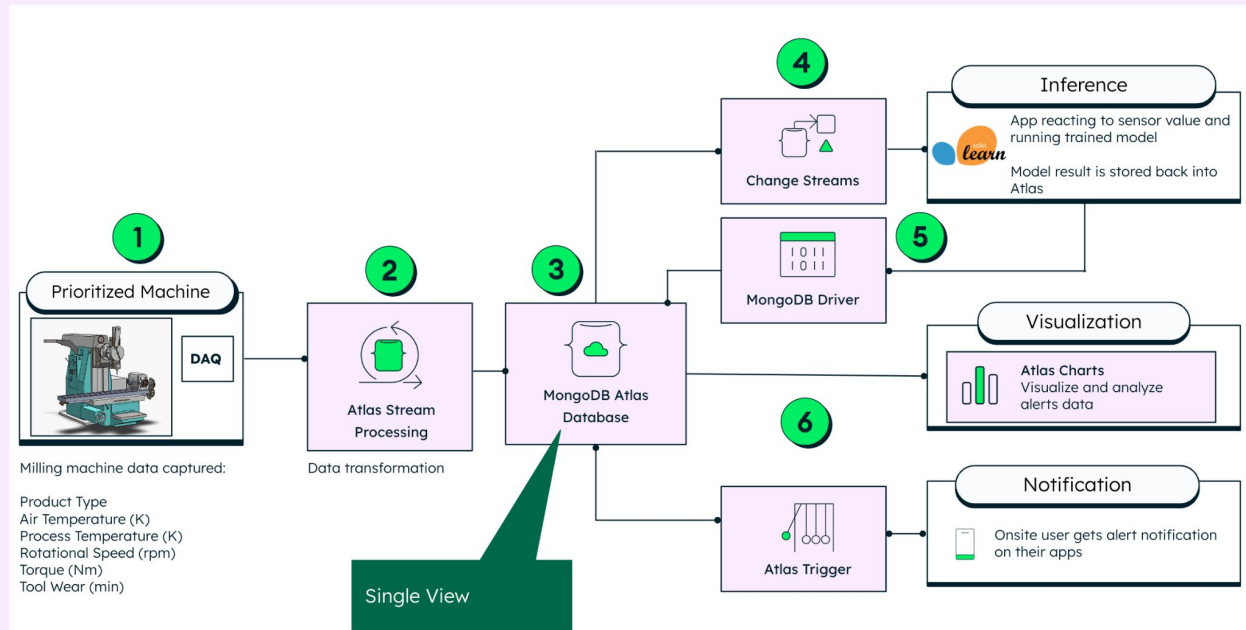


Figure 10: Smart milling machine uses real-time data to predict failures

As seen in the Figure above, we have our prioritized machine, which is a milling machine with attached sensors that collect data such as air temperature, rotational speed, torque, and tool wear. This data will be processed through Atlas Stream Processing, enabling the processing of streams of complex data using the same data model and Query API used in Atlas databases. [Atlas Stream Processing](#) enables developers to build aggregation pipelines to continuously operate on streaming data without the delays inherent to batch processing. Results can be continuously published to MongoDB Atlas or to a Kafka

topic. This allows data transformation and enrichment before it even lands in the database.

Once the data is in MongoDB, another application can react to sensor values and run a trained model designed to predict failures. The model results can be stored back into Atlas (between steps 4 and 5/Inference in the Figure above). These results can then be visualized using [Atlas Charts](#). Finally, Atlas Triggers and Functions can be used to push notifications to on-site users. This establishes an end-to-end system for failure prediction.

Stage 3: Repair plan generators

Having identified the nature of the equipment failures, the implementation of a comprehensive repair strategy becomes paramount. First, we have to generate a maintenance work order. This order should include repair instructions, spare parts needed,

schedule, and resource availability information. In this case, both structured and unstructured data are involved. The repair instructions will come from the machine manual. For this process, MongoDB Atlas acts as the operational data layer, seamlessly integrating structured and unstructured data.

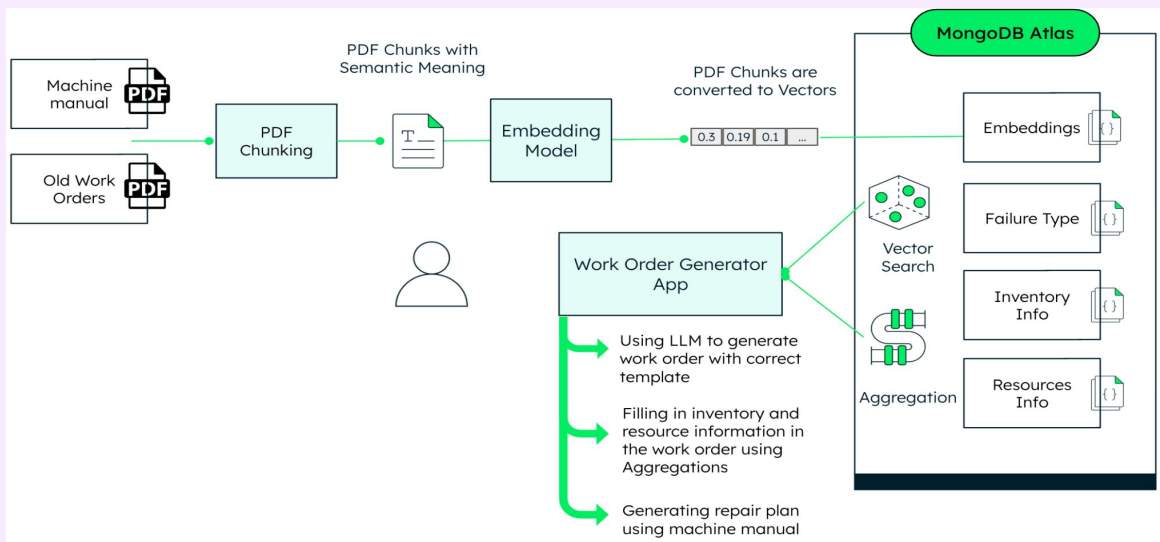


Figure 11: MongoDB Atlas as the operational data layer for structured and unstructured data

The Figure above shows the process of work order generation using generative AI. First, we must extract chunks of information from a milling machine's operating and repair manual, as well as from old work orders stored as PDF files, and convert them into vectors. These embeddings are then stored in MongoDB Atlas. MongoDB's versatility allows for the storage of both structured and unstructured data within the same database. Leveraging Atlas Vector Search and Aggregation pipelines, we can integrate this data to feed into a large language model (LLM) powering a work order generator application. The LLM analyzes the data to generate the appropriate work order and template, drawing from past examples. It populates inventory and resource details using aggregation techniques and structured data. Finally, it generates a repair plan similar to the

old work orders. What sets this approach apart is the ability to use the same MongoDB database to store structured data such as failure types, spare parts inventory, and resource information. By employing the aggregation framework to extract relevant information from structured data and vector search to glean insights from vectors, the LLM within the work order generator application gains contextual understanding.

This application seamlessly utilizes the LLM to generate work orders with the correct template, filling in inventory and resource details through aggregations, and ultimately creating repair plans based on machine manuals. This application can run inside a central maintenance management system.

Stage 4: Maintenance guidance generation

So we come to the last step: How can we use gen AI to enhance the operator or technician guidance to maintain the machine?

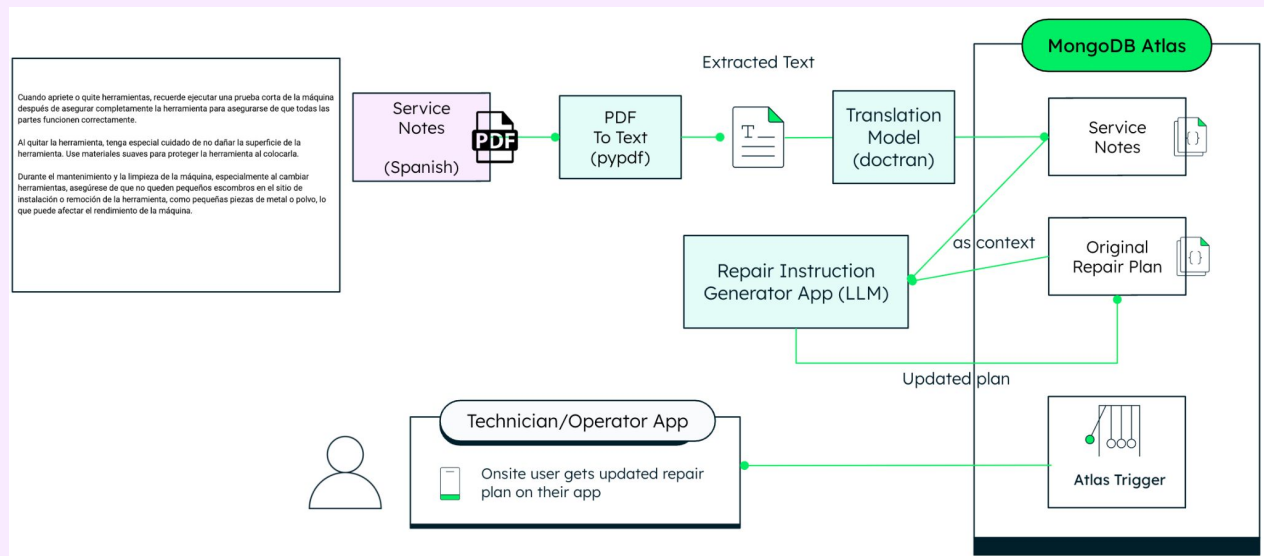


Figure 12: Using the RAG Approach for Operator Work Instructions

Let's walk through an example scenario here. The repair plan was generated in the last step. Now, the computerized maintenance management system (CMMS) has found some service notes uploaded to the platform by another technician, but they're written in another language, let's say Spanish. We can use the RAG architecture again to intelligently merge these service notes with the repair instructions generated in the previous step.

We first need to extract text from the PDF, translate it into English since our other data is in English, and then provide the service notes as well as the repair plan to the LLM as context. So, we have the original plan from the previous steps, and we combine it integrally using the LLM with the service notes obtained in this step. Note that we're not performing vector search here. Once the plan is updated,

then we can publish notifications down to the technician's application via Atlas Triggers and Functions.

In summary, we are essentially integrating AI and gen AI apps to implement an end-to-end predictive maintenance strategy (shown in the Figure on the next page).

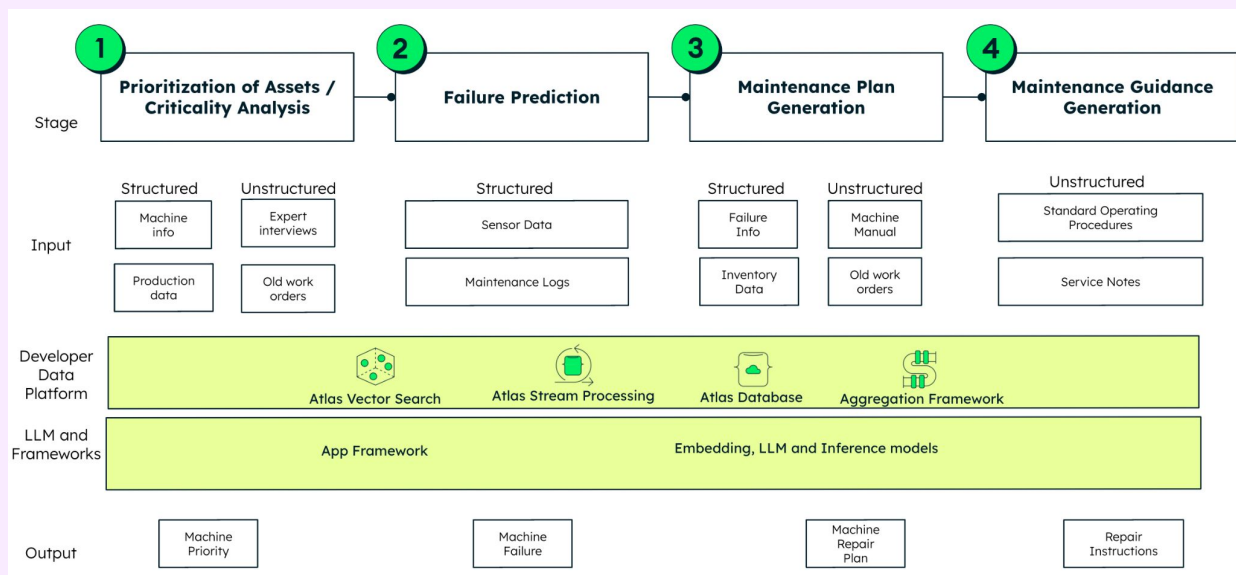


Figure 13: Model chaining with unified data store

Our input consisted of a combination of structured and unstructured data. We leveraged the various services offered by the MongoDB Atlas developer data platform, including Atlas Vector Search, Atlas Stream Processing, and, of course, the MongoDB database and aggregation framework. These features enabled us to provide the right context to the LLM and the appropriate data to the AI model.

Ultimately, we obtain the desired output at each stage, ranging from machine prioritization, failure type identification, and repair plan formulation, to instruction generation.

Solution Demo

Find out* how AI is being used in renewable energy by leveraging MongoDB Atlas Vector Search to drive efficiency through real-time, audio diagnostics.

*mongodb.com/solutions/solutions-library/real-time-audio-based-ai-diagnostics

Knowledge Management

The Unique Challenges of Preserving Knowledge

Preserving and maintaining knowledge in manufacturing is just as much a challenge as its accessibility. Due to aging population worldwide, as experienced workers retire, valuable tribal knowledge is lost. Transferring their expertise to the in-experienced workforce is difficult. Siloed data resulting from mergers and acquisitions or legacy systems makes it even harder to consolidate knowledge for decision making.

How Generative AI and Atlas Vector Search Help

Manufacturers can capture and index the valuable knowledge left by experienced

workers, including both textual and unstructured information. Creating semantic vectors from these documents, manuals, and notes simplifies this process. The task of locating and transferring knowledge from data silos gets eliminated by indexing and generating vectors from a wide range of data sources, encompassing both structured documents and unstructured data such as handwritten notes. This enables users to perform cross-system searches using natural language queries facilitating seamless access to information across different platforms.

Preserving Expertise through MongoDB Atlas

Documents and Manuals: Operation manuals, maintenance guides, process documents

Email Archives: Archived email communications related to specific projects, issues, or processes

Project Reports: Reports from completed projects, including success stories, lessons learned, and best practices

Standard Operating Procedures (SOPs): Any existing SOPs that experienced workers have followed and contributed to

Incident Reports and Case Studies: Records of past incidents, near-misses, and case studies

Meeting Notes and Minutes: Notes from meetings, particularly those involving experienced workers

Interviews and Personal Conversations: Transcripts or recordings of interviews with retiring workers, capturing their insights and experiences

Photos and Videos: Visual documentation of machinery, equipment, and processes

External Resources: Information from external sources, such as industry best practices

Technical Specifications and Blueprints: Detailed technical specifications, blueprints, and engineering drawings

Supplier and Vendor Documentation: Information provided by suppliers and vendors, including manuals and documentation

Regulatory and Compliance Documents: Documents related to industry regulations, safety guidelines, and compliance standards.

Customer Feedback and Quality Reports: Feedback from customers, quality assurance reports, and data on product defects or improvements

Employee Training Materials: Training materials used for onboarding and skill development

Embedding
Creation

MongoDB Atlas



Vectors, Core &
Metadata Store

Example LLM Prompts

- How to calibrate the de-palletizing robot?
- How to fix the shot peening machine?
- Give me summary of technical manual of my collaborative robot

Figure 14: Preserving knowledge in MongoDB Atlas

Knowledge collection from shop floor

The time being wasted for consolidating data from different systems to take decisions on a daily basis can be heavily reduced by using gen AI. A shift leader, as an example, spends a lot of time to collect data from different sources like MES, SCADA, or from handwritten notes of the night shift workers to get an overview about the condition of the equipment after the last shift. Traditionally it takes a lot of time to collect all that data from various sources and locations for getting a holistic overview to understand current safety, maintenance, inventory and quality needs.

A shifter leader may struggle to make timely decisions due to the fragmented nature of the data, leading to delays in addressing production issues and optimizing workflows.

With a knowledge management application on the shopfloor all the input from the production equipment as well as from the workers can be collected. The application takes all the structured and unstructured text input and categorizes it into one of many categories (defects, breakdowns, alarms, etc.).

Having all that data collected, contextualized, and indexed allows a chatbot application to get an immediate overview on the status of the shopfloor by prompting: Provide me a list of machines with problems in the last shift, followed by prompting on how to solve that problem. The RAG application can therefore use all the preserved information from the experienced workers, stored as vectors in MongoDB Atlas.

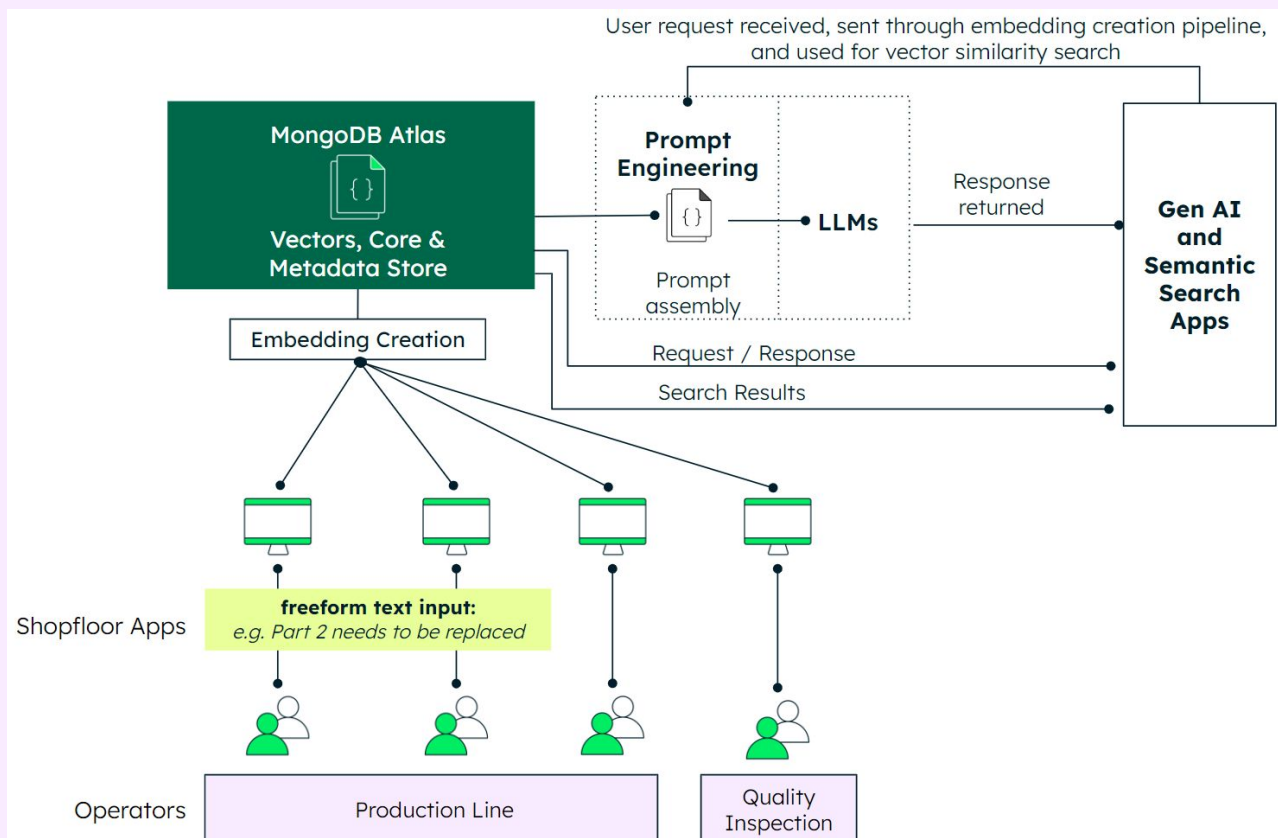


Figure 15: Knowledge management architecture



Eni makes terabytes of subsurface unstructured data actionable with MongoDB Atlas

Based in Italy, **Eni** is a leading integrated energy company with more than 30,000 employees across 69 countries. Its operations vary from exploring and drilling for natural gas and oil to cogenerating electricity, renewables, biorefining, and chemical production.

Eni partnered with **MongoDB Consulting** for training and to support the migration of workloads into **MongoDB Atlas**. Eni wanted to move to a managed service with a seamless user experience and easy-to-use interface for developers.



Improving staff productivity at Enel using Amazon Bedrock

Enel is a leading integrated electric utility with a presence across 32 countries and an 82-GW generation capacity

Enel identified the opportunity to use generative AI to boost IT service desk efficiency by extending automation to nontrivial tasks through basic troubleshooting, providing resolution steps and ticket routing without human involvement.

The solution is designed around a retrieval-augmented generation architecture using Amazon Bedrock.

With MongoDB Atlas, Eni users can quickly find data spanning multiple years and geographies to identify trends and analyze models. **MongoDB Atlas Search** also assists by filtering out irrelevant documents. The team also integrated AI and machine learning models with the platform to make it even easier to identify patterns.

“MongoDB Atlas isn’t just a database, it’s a complete set of products and services. It’s cloud agnostic and combines rich functionality with the flexibility we needed to make it our own.”

Learn more*

Sabato Severino

*Senior AI Solution Architect for
Geoscience at Eni*

*mongodb.com/solutions/customer-case-studies/eni

The solution uses Amazon Titan, a family of models exclusive to Amazon Bedrock. Specifically, it uses the Amazon Titan Text Embeddings model to generate embeddings (vectors capturing semantics of text) from Enel’s knowledge base, which consists of a series of runbooks containing incidents classes, preconditions, root causes, resolutions steps, and operations information related to the applications. Embeddings are computed and persisted in a vector database instance using **MongoDB Atlas Vector Search**, which supports similarity search.

Learn more*

*aws.amazon.com/blogs/industries/improving-staff-productivity-at-enel-using-amazon-bedrock

Other Notable Use Cases



AI plays a critical role in fulfilling the promise of Industry 4.0. There are numerous other use cases of AI that can be enabled by MongoDB Atlas.

Logistics Optimization

AI can help optimize routes resulting in reduced delays and enhanced efficiency in day-to-day delivery operations.

Quality Control and Defect Detection

Computer or machine vision can be used to identify irregularities in the products as they are manufactured. This ensures that product standards are met with precision.

Production Optimization

By analyzing time series data from sensors installed on production lines, waste can be identified and reduced, thereby improving throughput and efficiency.

Smart After Sales Support

Manufacturers can utilize AI-driven chatbots and predictive analytics to offer proactive maintenance, troubleshooting, and personalized assistance to customers.

Personalized Product Recommendations

AI can be used to analyze user behavior and preferences to deliver personalized product recommendations via a mobile or a web app, enhancing customer satisfaction and driving sales.

FOR MORE INFORMATION AND RESOURCES

Visit **MongoDB Atlas** for **Manufacturing and Motion***

*mongodb.com/solutions/industries/manufacturing

Contact Information



Dr. Humza Akhtar

Manufacturing & Motion
Industry Solutions Principal
humza.akhtar@mongodb.com



Raphael Schor

Manufacturing & Motion
Industry Solutions Principal
raphael.schor@mongodb.com