/ WHITE PAPER

CRATE.IO

Industrial IoT in manufacturing: why the industry needs a specialized database

A new industrial revolution: the era of IIoT

The manufacturing industry is going through a deep transformation. The days of the simple automatization of processes are over: in the factories of the future, decision making will be ruled by data.

For decades, all kinds of on-site monitoring devices have been collecting huge streams of data during industrial operations, but this data has remained mostly unused. Traditionally, the different signals have been monitored individually, and the action taken locally; this is about to change. Innovation will be a key factor to survive in the future marketplace, and to unlock the potential of machine data is the next step for the industry.

Big data analysis, artificial intelligence, and cloud computing are changing the world as we know it. Bringing the power of the Internet of Things (IoT) to the industrial world, the adoption of Industrial IoT (IIoT) systems in factory floors is leading to a new industrial revolution, and a new generation of *smart* factories are being born.

In the modern industry, processes will no longer be treated in isolation; data will not only be collected, but processed, stored, exchanged, and analyzed. The term Industry 4.0 is applied to this new generation of digitalized factories, where the proper collection and processing of data makes it possible to remotely monitor operations in real-time, allowing to take action immediately when a failure occurs. In fact, failures will often be identified before they happen, using predictive maintenance and machine learning algorithms.

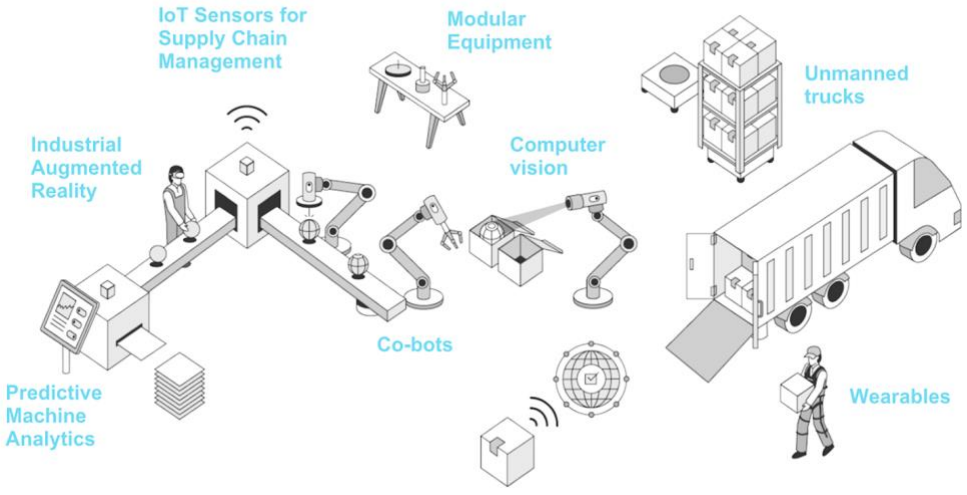

The factory of the future will be treated as a whole, taking advantage of modern technologies

The potential of IIoT is enormous: de-centralized operations, increased efficiency and safety, reduced labor costs, more effective business strategies, and easy visualizations of the complete operations. All is possible by using the power of data.

However, to effectively put IIoT into practice is being a challenge for manufacturers. With about half the IIoT projects failing in the proof-of-concept stage, industries are facing several difficulties, including the lack of training of the existing personnel and a resistant corporate culture.

Changing paradigms is never easy. Nonetheless, there is a factor commonly found among failed IIoT projects: a lack of understanding about the technological needs of IIoT. Many projects have failed due to the use of inadequate tools, making mistakes that would have been easily avoidable.

Actually, to successfully implement IIoT practices is not as intimidating as it might seem, but in order to succeed it is key to use the right technology. Since data management is the spinal cord of IIoT processes, this is especially relevant in what respects to the database.

New challenges, old solutions? Choosing the right database

Handling IIoT workloads: requirements

From the point of view of the database, the IIoT is a use case with a very specific combination of needs:

- Industrial processes generate huge volumes of data. In terms of data volume, IoT applications had posed entirely new challenges for the databases, and the degree of difficulty increases even more in the case of industrial IoT. During industrial processes, different types of signals and sensors are being collected continuously: the database not only needs to be able to handle large data volumes, but to allow the querying of a wide variety of data types.
- In order to allow the monitoring and control of operations, the real-time processing of data is key. This implies that the database for IIoT must be able to handle complex, highly concurrent queries in milliseconds. At the same time, the database must allow the analysis of historic data for using machine learning algorithms.
- The needs of the industry are dynamic: new factories are open, the volume grows, and new processes are established. To perform expensive, tricky database migrations every time this happens is not acceptable for the future industry. The database for the IIoT era must be dynamic, ready to be adapted and optimized.
- In the industry, all processes need to be done in the most efficient way possible. Of course, the database is not an exception. This is a crucial differencing factor of IIoT applications, especially those applied in manufacturing industries with high production volumes and low margins. To choose a database operating with cost-efficiency in mind is a necessity, to avoid unnecessarily high database costs in the future.

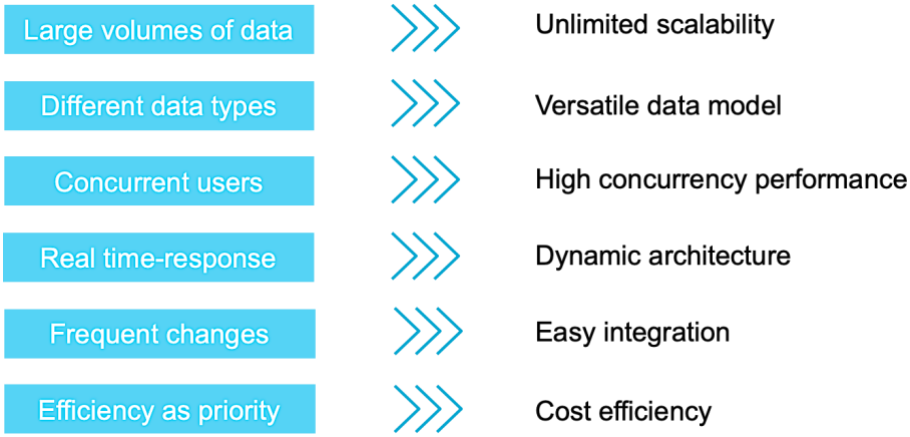

The needs of the IIoT use case determine the requirements for the database

Scalability, performance, efficiency: a hard-to-find combination

All the characteristics presented above are crucial for the success of IIoT. However, it is not easy to find a database with such a combination of properties. Databases are not a one-for-all solution: every use case comes with a set of priorities, and for each application, some databases will naturally be a better fit than others. IIoT is not an exception: being a relatively new use case, most of the existing database solutions don't quite match the IIoT needs.

For example, popular databases among industrial organizations, like SQL Server and Oracle, cannot handle the complexity and volume of IIoT. They are expensive to scale, and usually unable to process the huge data volumes involved. At the same time, they are not built to handle the query complexity involved in modern industrial applications, and their real-time capabilities are not powerful enough to deliver millisecond responses in these cases.

It may seem that these issues would be solved by migrating to a NoSQL database like MongoDB or Cassandra, and indeed they offer significant advantages. However, they also bring new problems to the table. NoSQL databases offer excellent scalability, and their distributed architectures allow them to perform complex, flexible queries. But they also imply intricate infrastructures that require intensive planning and administration, together with specialized engineers.

Besides, NoSQL databases are usually difficult to adapt or to integrate with other tools, and they come with high costs in terms of hosting, integration and administration. In addition, their speed often degrades rapidly when querying data over weeks or months. The ability to process longer queries is important for industrial IoT applications, since it allows to perform predictive analysis with accuracy.

So what about time series databases? Time series data are a big part of the data generated by industrial processes, so it is understandable that time series databases might seem like a good candidate for the modern industry. However, a time series database should not be foundational to an IIoT system.

The time series data typical of IIoT is generated by a huge number of sensors, involving highly concurrent clients and different data types. This scenario differs substantially from the design window of most databases specialized in time series, whose functionality and scalability typically suffer from intense parallel usage.

IIoT requirements	Traditional RDBMS	NoSQL	Time series databases	CrateDB
High data ingestion	No	Yes	Yes	Yes
High data versatility	No	Yes	No	Yes
Real-time performance	No	No	Yes	Yes
Simple scalability	No	Yes	No	Yes
SQL access	Yes	No	No	Yes

Besides, time series databases are optimized for the exclusive handling of time series data: by themselves, they are not able to process the unstructured data that is also a part of IIoT operations. Manufacturers might need to run two databases in parallel, which results in complex setups, large cloud footprints, querying disadvantages, and higher costs.

How to make it work

CrateDB: a database built for the modern industry

Existing databases don't fit the needs of modern industrial applications: CrateDB was created to fill that void.

It is a database purpose-built for the modern industry, designed to offer the combination of features that are most important for the success of IIoT.

Combining the best of both worlds, CrateDB is a distributed SQL database built on top of a NoSQL foundation. It specializes in cases where huge volumes of data are continuously coming in from different sources, providing real-time responses and allowing to visualize, monitor and control complex operations in real-time. And it all is done while keeping the costs low.

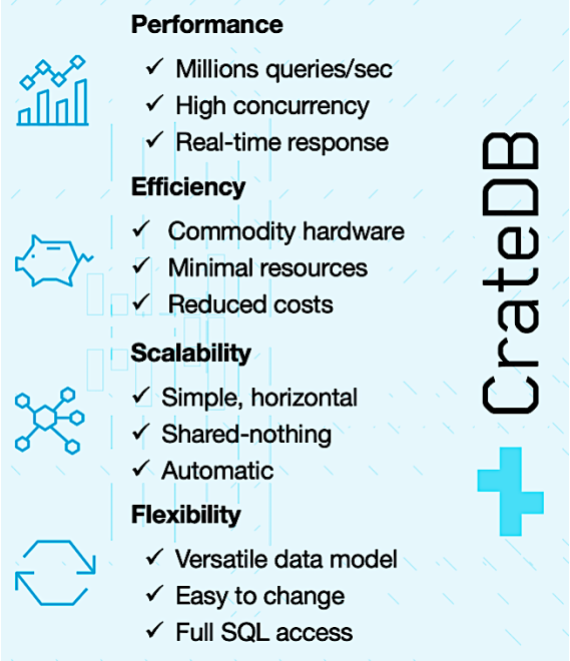

The future of manufacturing, today: a real story

ALPA, the billion-dollar plastics manufacturer, has succeeded where others have failed. Being the world leader in the development and production of plastic packaging solutions, and with about 200 production plants spread across 46 countries, ALPLA is an example of what can be achieved by adopting IIoT practices in manufacturing, and of how it is possible to succeed if the right tools are used.

After implementing the first phase of their digital transformation using CrateDB Platform, ALPLA's production plants improved their NEE (Net Equipment Efficiency) by above 90%, and manual process adjustments were reduced by 50%. The result was a six-figure saving in 2018.

This is how it was done.

Motivation

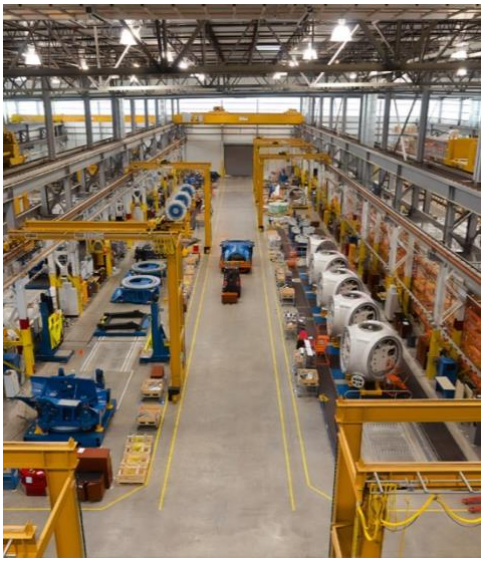

APLA began its IIoT transition to achieve data-driven manufacturing. They realized that **t**o identify and correct issues within their factories was a slow and labor-intensive operation: it required multiple workers to be constantly moving around checking up the machinery, needing to cover considerable distances. This meant that many processes were often affected before the operator had the opportunity to spot the error or to correct it, with the corresponding loss of valuable machine time and with a considerable amount of raw material wasted.

Additionally, ALPLA was growing. As more and more ALPLA plants were opening around the world, the challenge of hiring and training enough specialized personnel increased. Modern manufacturing machines are often complex, so every plant required to have trained experts present on every shift: to find such a high number of specialized workers was becoming increasingly difficult, and the related costs were high.

The number of factories kept increasing as the business grew. A plastic bottle is a mass-produced, cheap, light product, so to transport it long distances it is extremely inefficient economically: the solution lies in building production facilities close to the customer, meaning that the manufacturer needs to operate, control, and regulate a large number of production plants.

The initial struggle

Initially, ALPLA tried to keep using the data system they already had in place, SQL Server. However, it was unable to cope with their data requirements. To achieve a real-time response was critical for APLA, in order to implement a system that allowed workers to address possible failures as quickly as possible; but as a result of ALPLA's rapid growth, the database performance was significantly degraded.

The data requirements were severe: 900 different sensor data structures, and thousands of readings per second generated by each production line. However, when ALPLA switched to data technologies specifically built for IIoT, the improvement was immediate. When CrateDB was implemented, data queries start running 250 times faster than with SQL Server.

The result

After implementing CrateDB Platform, the days of the red-light alarm are over. Instead of regularly checking the machinery looking for failures, the production lines are now checked automatically on a 24/7 basis, monitoring aspects like quality control, alarms, processes, parameters, and procedural alerts.

These checks are pre-defined for a specific production process, being easily scalable to multiple plants. This allows ALPLA to implement the model quickly in new factories, without the costs related to duplicated work.

When a failure is detected by one of the checks, a note is automatically sent to an employee on the factory floor, who will hear a message delivered through Bluetooth earbuds. If the issue detected requires the assessment of an expert, an alert is sent to the Mission Control Room.

The Mission Control Room allows ALPLA to monitor the operations of multiple plants with only one set of experts, reducing substantially the labor costs related to specialized personnel. By the use of centralized dashboards, the expert operators have insight into the current state of all the processes and machinery of multiple factory floors. At the same time, they have access to analytics from related historical data. From this high-level vantage point, operators make informed decisions, notifying the staff on the production floor if action is necessary.

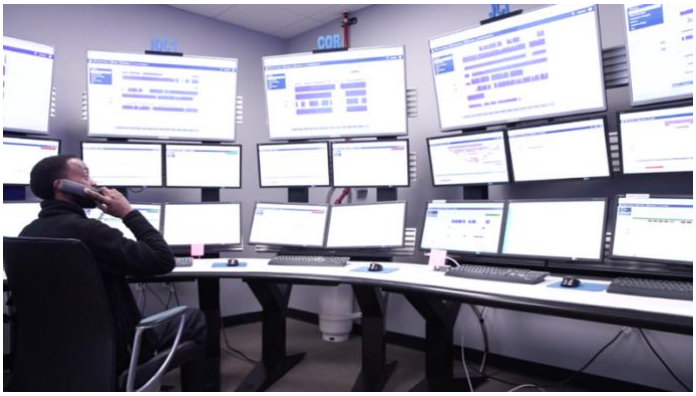

In the Mission Control room in McDonough, GA, operators monitor and support the operation of 11 plants across the USA

Conclusions

There's no doubt: IIoT is the future of manufacturing. The digitalization of industrial processes has the power to lead incredible long-term profits: it increases the efficiency of the operations, the degree of automatization, it reduces waste, and it allows to set up new factories in areas with scarcity of specialized personnel. Besides, the access to historical data opens the door to AI and machine learning, further improving the efficiency of the operations by predicting future failures and quickly correcting them. The success of the ALPLA story speaks for the transformative power of the IIoT, and it is only the beginning.

However, in order for IIoT projects to succeed, it is crucial to use technologies that have been specifically optimized for this use case. The IIoT implies considerable complexities and particularities in terms of data, and it is unlikely that a database not specialized in IIoT will provide all the necessary requirements.

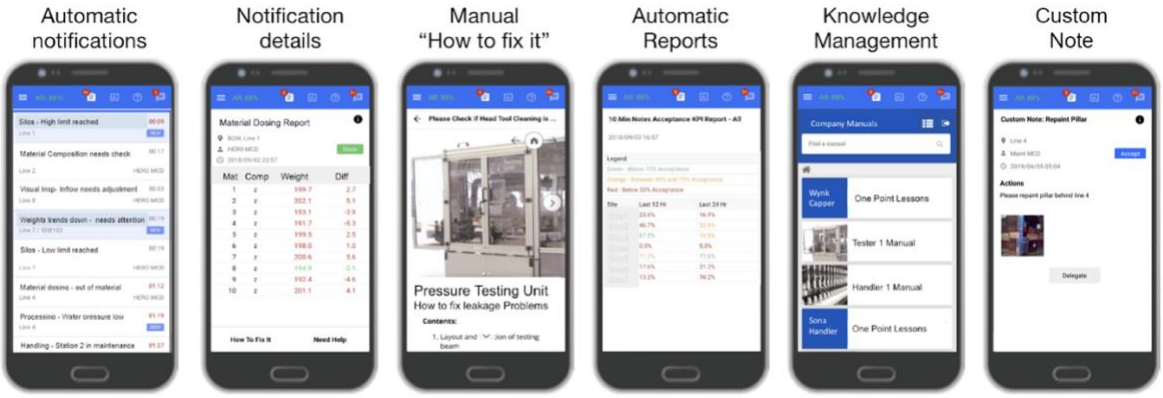

The Notes App

This problem is solved with CrateDB, a database built for the modern industry. It provides excellent real-time performance for IIoT workloads, simple scalability, and easy accessibility and integration. Besides, CrateDB is designed to be as efficient as possible, both in terms of costs and resources.

Please, reach out to us if you have any questions about CrateDB or your use case. We'll be happy to help.

If you want to keep learning about the Industry 4.0 and IIoT, check out the materials on our resource library!