



Real-Time Monitoring: Selecting the right database

E-BOOK

 **CrateDB**



In today's fast-paced business world, real-time monitoring has become essential for companies of almost all sizes and industries that want to stay competitive.

Continuously tracking and analyzing data as it is generated enables businesses to identify anomalies, prevent issues before they arise, and make informed decisions.

Whether it's monitoring equipment performance on a factory floor, tracking customer behavior on an eCommerce site, or analyzing financial market data, real-time monitoring provides businesses with powerful insights.



However, to fully leverage the benefits of real-time monitoring, businesses need the right infrastructure in place to support the large volume of data.

In this ebook, we will explore the key benefits of real-time monitoring, the critical infrastructure components required to implement it effectively, and why CrateDB is the database of choice for real-time monitoring projects.

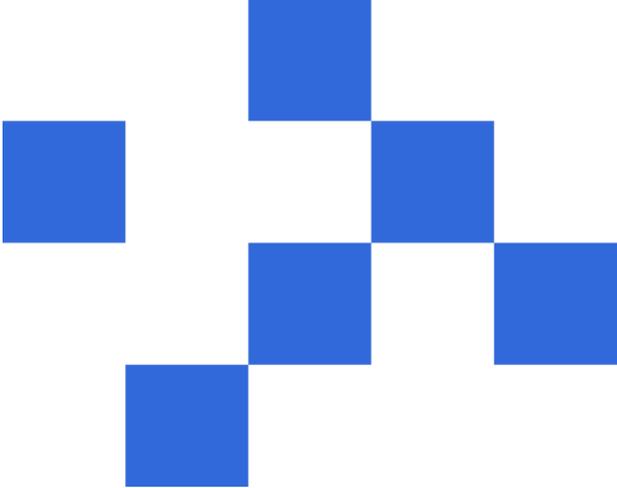


Table of contents

1. What is real-time monitoring?
2. Why is real-time monitoring important?
3. How different industries leverage real-time monitoring
4. Choosing the right infrastructure for real-time monitoring
5. Why CrateDB is the ideal database for real-time monitoring
6. Powering real-time monitoring with CrateDB: 3 real-world examples
7. How can your business take advantage of real-time monitoring?

What is real-time monitoring?

Real-time monitoring uses technology to collect and stream data, usually via sensors or connected devices, with practically zero latency between collection and delivery.

This provides access to business-critical information and trends in real-time, enabling businesses to make better-informed decisions and take proactive action.

Real-time monitoring connects all the dots with continuous data streaming that provides greater visibility into processes like production and shipping.

Why is real-time monitoring important?

Real-time monitoring has a demonstrable impact on business success. A recent study found that **62%*** of businesses see an improvement in process efficiency after implementing real-time monitoring, and **80%*** of them report increased revenue.

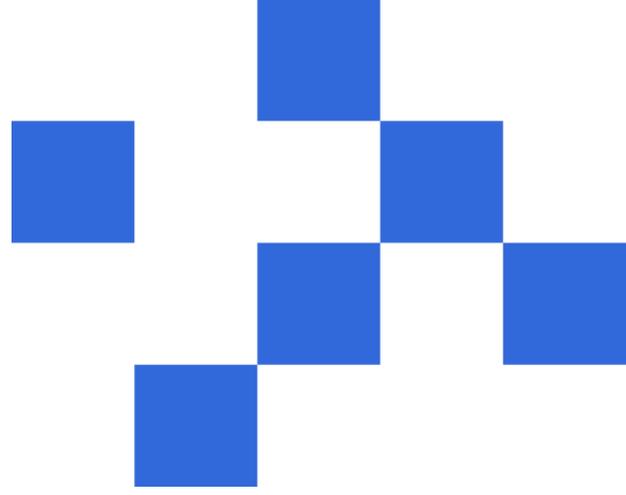
Why exactly is that?

Firstly, the main benefit of real-time monitoring is the ability to gain actionable insights in real-time. This means that businesses can identify potential issues and take corrective action before they have a negative impact on operations.

For example, a manufacturing plant could use real-time monitoring to track the performance of their machines, detect any abnormalities and perform proactive maintenance before a breakdown occurs, avoiding costly downtime.

Another clear use case for real-time monitoring is eliminating human errors.

*Source: CISION [here](#) and [here](#)



Another clear use case for real-time monitoring is eliminating human errors. By automating manual checks that are prone to mistakes, businesses can ensure greater levels of consistency and accuracy. Real-time monitoring also increases transparency, enabling businesses to keep customers, partners, and key stakeholders up-to-date throughout shipping and other key processes.

Key benefits of real-time monitoring:

- Get actionable insights in real-time
- Reduce the risk of human error
- Increase transparency
- Save costs and boost profit



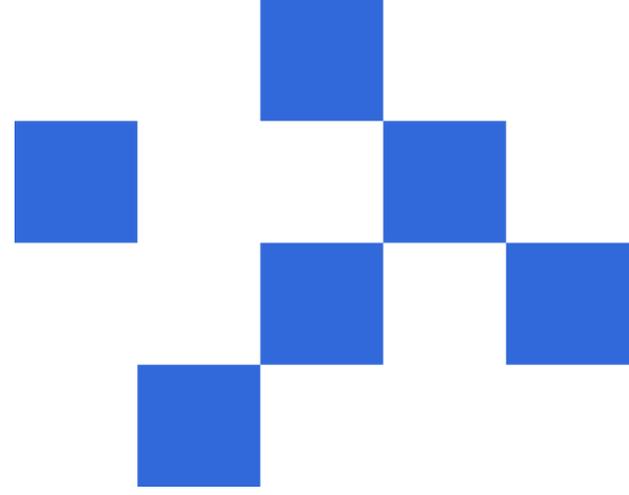
How different industries leverage real-time monitoring

Manufacturing

Real-time monitoring enables manufacturers to keep a close eye on their production lines and identify potential issues before they cause significant problems. For example, sensors placed on machines give valuable information about their performance, and if any issues arise, the manufacturer can quickly address the problem before it becomes a bigger issue.

Healthcare

Real-time monitoring shows significant potential in the healthcare industry. Using wearable devices to monitor patients' vital signs enables doctors and medical staff to provide timely care, identify emergencies, and alert family members when needed. Real-time monitoring can even be used to improve cancer treatment, whereby data is collected from patients and shared with oncologists to enable remote checkups and more reliable tracking of changes in symptoms.

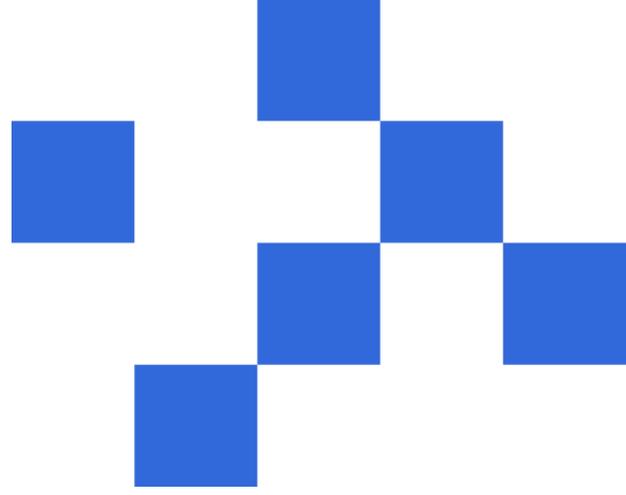


Transportation & Logistics

Real-time monitoring, usually via GPS tracking, can help optimize transportation routes and improve overall efficiency for logistics companies. Tracking the location of goods and vehicles, as well as adjusting routes in real-time based on traffic conditions, helps ensure that they arrive at their destination on time.

Energy

Real-time monitoring plays an essential role in the energy sector as it enables power plants and energy facilities to optimize their operations. This is especially important for renewable energy systems, such as wind and solar power, which are exposed to fast-changing conditions. Sensors and monitoring systems can track the performance of the turbines and panels, monitor weather conditions, and adjust the operation of the system to maximize efficiency and output.



Agriculture

So-called smart farming is another key use case for real-time monitoring, with the global market value of smart agriculture expected to reach \$34 billion by 2026. Monitoring weather, soil and other data enables farmers to optimize crop yields, plan optimal planting schedules, and reduce waste. For instance, sensors placed in fields can monitor soil moisture levels and adjust irrigation accordingly, and real-time weather data helps farmers decide when to plant and harvest crops.

Choosing the right infrastructure for real-time monitoring

Product
Product_id
Product_name
Amount
Price

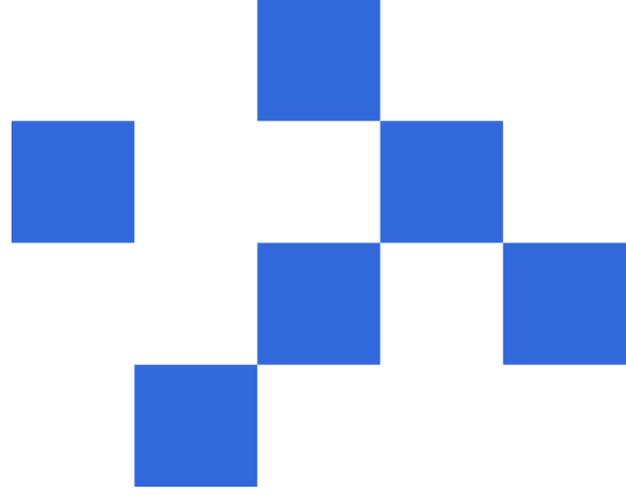
Invoice
Invoice_id
Customer_id
Order_id
Product_id



Implementing real-time monitoring requires a solid infrastructure that is capable of processing and analyzing large amounts of data in real-time. The specific infrastructure required will depend on the type and volume of data being monitored, as well as the specific needs of the business. However, here are some key components that businesses typically need to support real-time monitoring:

Data collection:

The first step in real-time monitoring is to collect data from various sources, such as sensors, IoT devices, or APIs. The data needs to be collected and stored in a central location that is accessible to the monitoring tools.

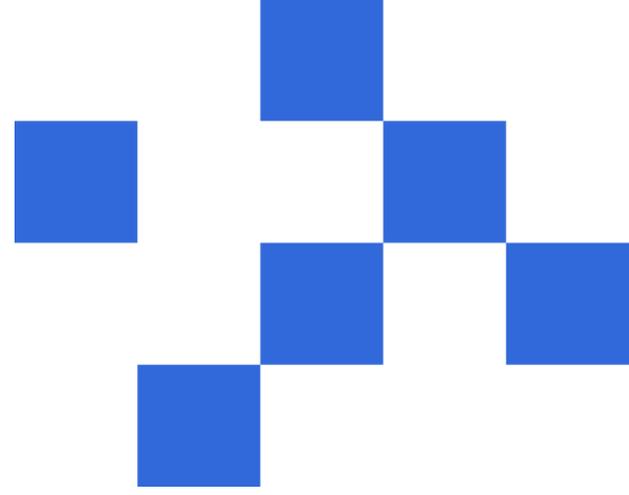


Data processing:

Once the data is collected, it needs to be processed to make it usable for real-time monitoring. This involves transforming the data into a structured format that can be analyzed and acted upon.

Real-time analytics:

To monitor data in real-time, businesses need to use analytics tools that can process and analyze data as it is collected. These tools can identify patterns, anomalies, and trends in the data, and provide alerts or notifications when certain thresholds are reached.



Visualization and reporting:

To make real-time monitoring actionable, businesses need to be able to visualize and report on the data in a way that is easy to understand. This can involve creating dashboards or reports that display real-time data in a clear and concise format.

Real-time monitoring typically generates huge amounts of data, so businesses need a scalable infrastructure that can handle large volumes of data without slowing down or crashing. Furthermore, businesses need to ensure that their systems are secure and protected from potential cyber threats.

By investing in the right infrastructure to support real-time monitoring, businesses can gain valuable insights into their operations and make data-driven decisions that can lead to improved efficiency, profitability, and customer satisfaction.

Console

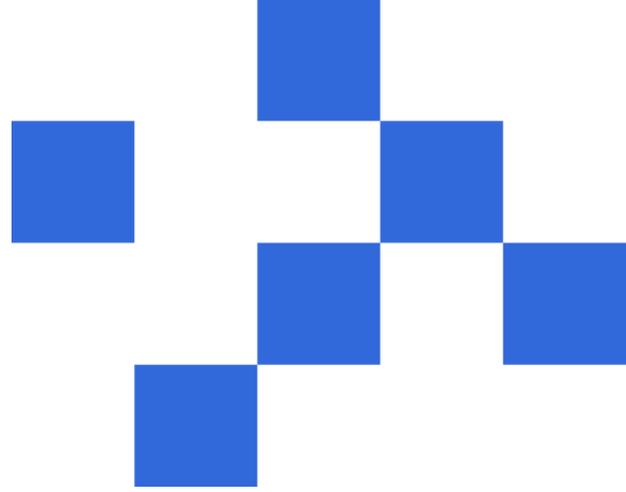
Why CrateDB is the ideal database for real-time monitoring



Real-time monitoring is highly valuable, but to get the most out of it, you need the right infrastructure to support it. Quick retrieval of data and real-time analysis are essential if the system is to handle large amounts of incoming data.

A traditional relational database may struggle with the sheer volume of data, and the high read/write requirements, leading to slow and unreliable performance.

That's where a purpose-built database is a real business advantage. CrateDB is a distributed SQL database that makes it simple to store and analyze large amounts of data in real-time. Designed for scale and high-speed retrieval, CrateDB delivers speed, reliability, and the power to process vast amounts data in real-time.



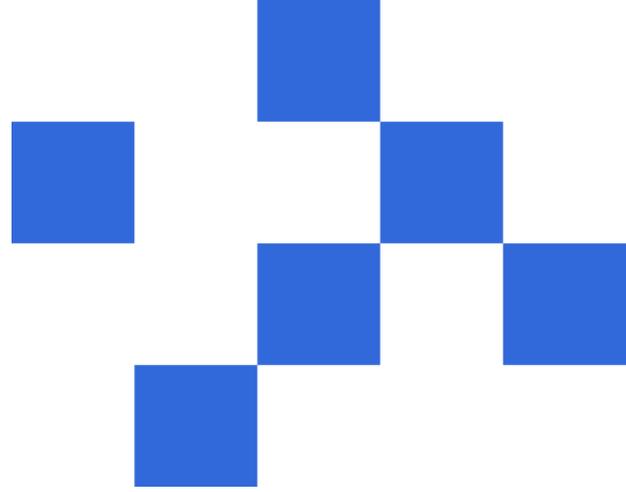
Here are some real-life stats from companies
CrateDB works with:

Data flowing from 100+
factories and 1000+
production lines

1,000 different sensor
message data
structures

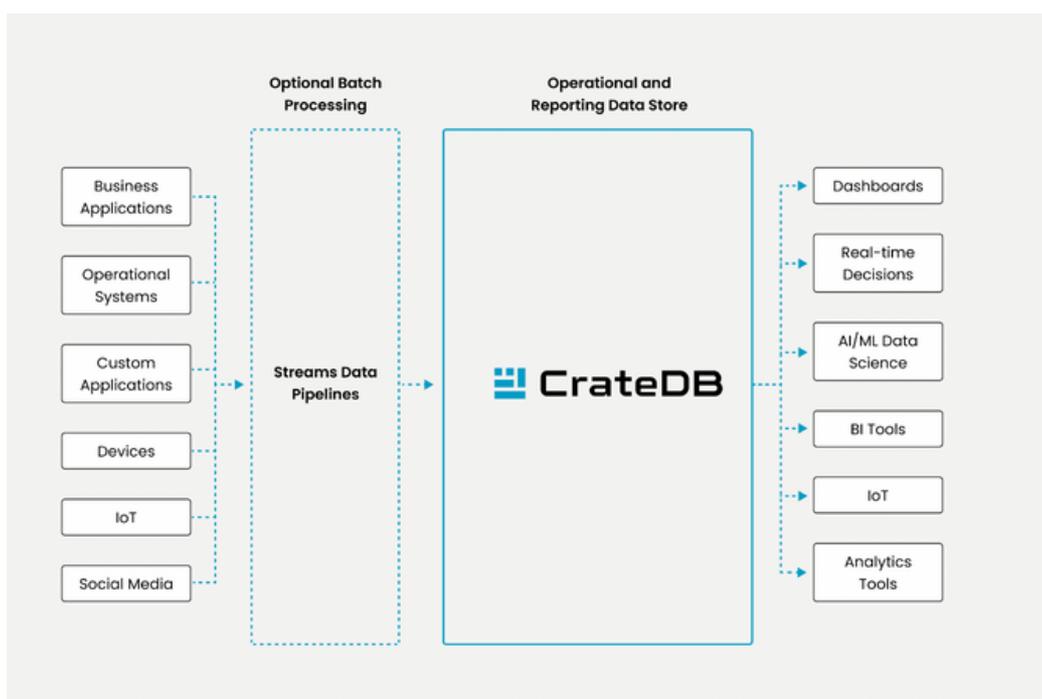
900,000 inserts per day
and 100+ terabytes of
retained data

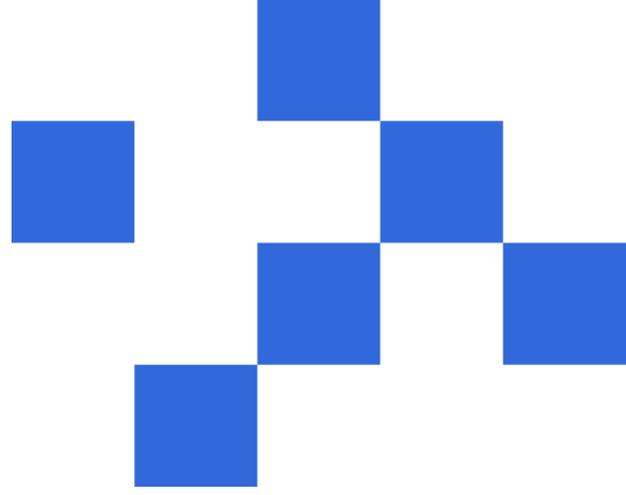
Real-time dashboards
executing ~ 1 million
queries per day



As CrateDB is highly scalable, it can handle data workloads of any size with ease. With a truly distributed architecture, CrateDB serves high ingest loads while maintaining sub-second reading performance in many scenarios.

As an ANSI SQL DBMS, CrateDB is easy to learn, easy to integrate, and does not lock you into proprietary data access interfaces. Here are some more reasons people choose CrateDB for real-time monitoring projects.





Process millions of data points per second

CrateDB can ingest, store, and process millions of data points per second thanks to its distributed processing, data partitioning, multithreaded design, and shared-nothing distributed architecture with masterless clustering.

Real-time query performance

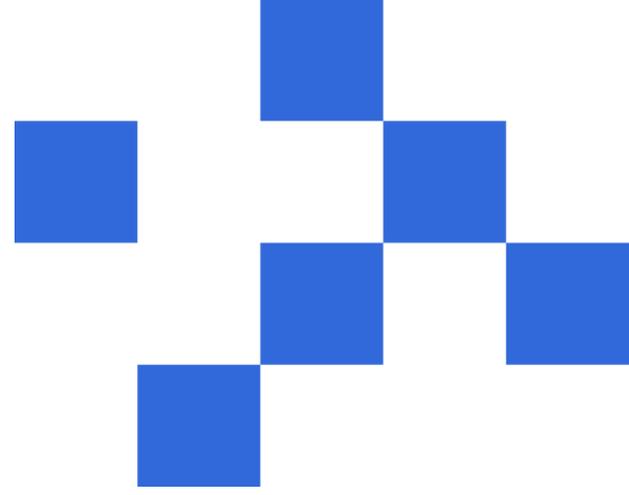
In-memory columnar indexes, field caches, and a distributed query execution engine can parallelize complex queries across the whole cluster.

Powerful analytics in real-time

Partitioning, parallel processing and in-memory columnar indexes enable real-time, complex analytics, not just simple aggregates.

Extensible time series data models

With elegant JSON handling and Dynamic Schema, CrateDB automatically adapts schema to new JSON structures.



Time series data & SQL interfaces

Built-in interfaces enable you to tap into Azure IoT Hub, Azure Event Hubs, Prometheus, Telegraf, and almost any SQL tool – such as Grafana, JDBC, and REST – via Postgres wire protocol.

Scalability

Elastic scaling, partitioning, sharding, and replication power fast, always-on performance that stays consistent as data volume and concurrent clients grow.

Built-in high availability

Automatic replication and self-healing ensure ultimate reliability and non-stop performance.

Simple deployment & easy onboarding

You can deploy CrateDB with container tools like Docker and Kubernetes and scale up or down with a single command to linearly increase or decrease capacity. Plus, as CrateDB is an SQL database, developers can quickly get up to speed.

Powering real-time monitoring with CrateDB: 3 real-world examples

Let's dive into two examples of how companies are using CrateDB for real-time monitoring.



Founded in 2017, O-CELL is a French company that aims to reduce the environmental footprint, creating a solution to reduce energy and water waste and lower the ecological impact of infrastructures.



Bitmovin is a leading provider of streaming solutions and video analytics. Founded in 2013 in Klagenfurt, Austria, the fast-growing company needed a high-performing database that could handle real-time analytics on large and fast-moving datasets, at scale.



SPGo! is part of PETROMIN and uses IoT-based predictive maintenance to help businesses improve the efficiency and profitability of their mining operations.

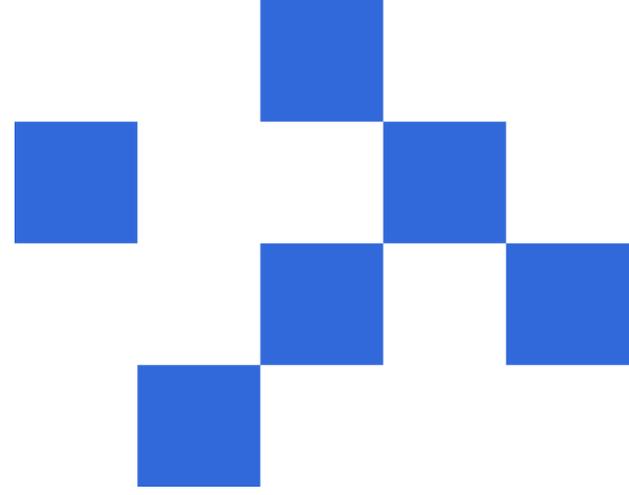


The challenge

O-CELL designed a flexible solution built around dedicated services to ingest, process, and present the data. This multi-purpose platform collects near real-time energy consumption prediction and other environmental parameters and provides easy-to-read dashboards for the city's representatives, and access this data for the technical staff and detect abnormal behavior.

The solution

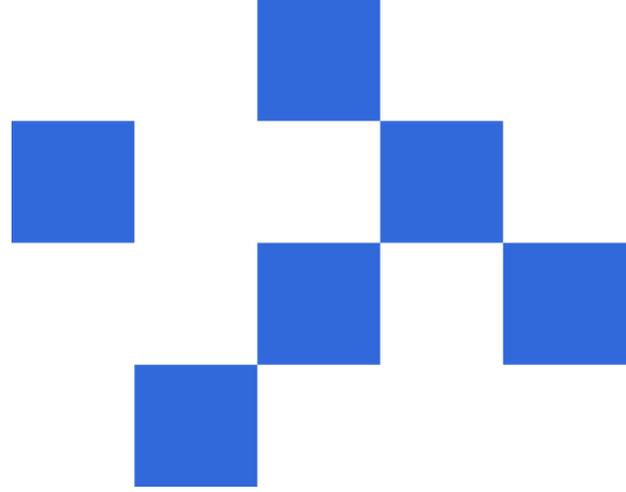
O-CELL has been using CrateDB since the early stages of its project. With the system's evolution over time, CrateDB has supported O-CELL in every phase of the business. As a small company, they have limited resources and must use them wisely. With a database like CrateDB, which needs very little maintenance, their team doesn't waste time-solving, for example, replication issues.



Charles-Edouard Ruault, O-CELL's Co-Founder, explained how they designed the data schema in the database, which stored the data collected by the sensors. One of the basic goals of the system was to be extensive and able to store many data types and add new ones easily.

The data is sent from the sensor as binary frames and encoded according to the sensor's manufacturer's specifications. For example, the IoT endpoint service received this frame storage in a dedicated table with made that data reflecting the signal quality. The decoder service will then fetch those frames, process them using the corresponding algorithm, normalize the data, and store all the data as indicators. So these indicators are just time series that can contain different data types sent from the same sensor, and at the same time.

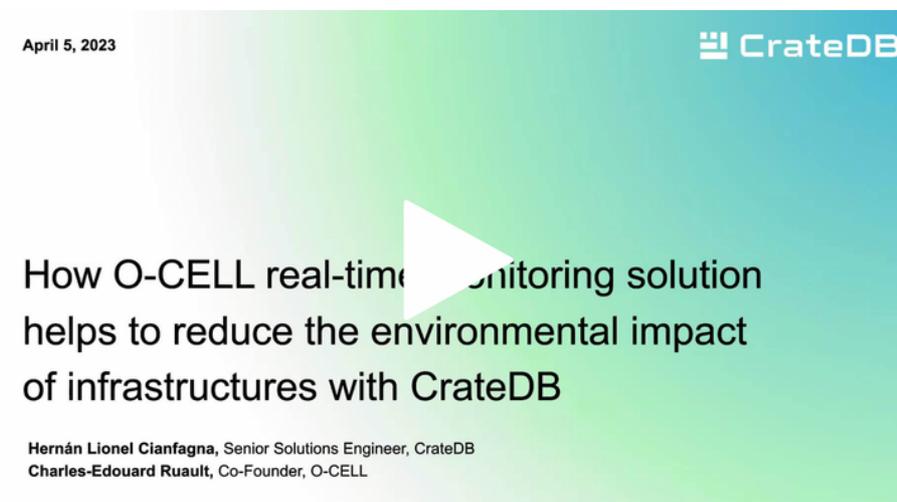
CrateDB allows O-CELL to store all the indicators of produce by the sensor simultaneously into one single database row, which is excellent for storage and query performance.



We've been pretty happy with this setting, and the system is running smoothly.

Charles-Edouard Ruault, Co Founder, O-CELL

Watch the [45-minute webinar](#) with O-CELL to learn more about their use case or check the summary blog post to quickly check the webinar highlights.



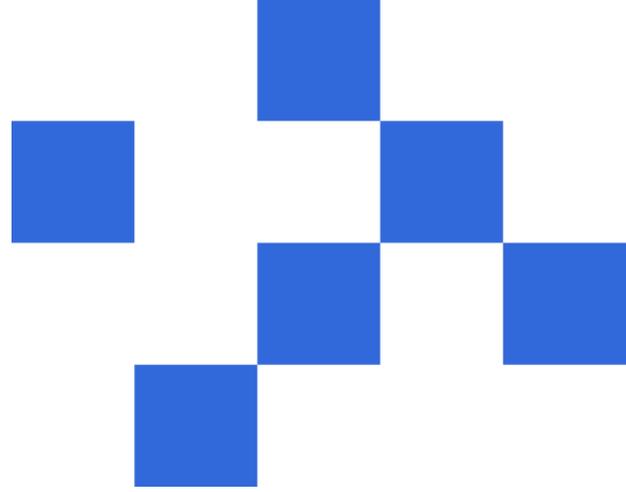


The challenge

When Bitmovin launched their analytics platform, they knew they needed powerful infrastructure to support it, so they tested different databases and APIs. They were looking for a highly reliable, scalable database with built-in fault tolerance and end-to-end latency of 5-10 seconds. Based on their initial tests, CrateDB was the clear winner thanks to its scalability, reliability, and ease of use. Onboarding was easy: CrateDB uses SQL which was immediately familiar to Bitmovin's developers.

The solution

Today, Bitmovin produces billions of rows of data related to video playback experiences and stores this data in CrateDB for real-time analytics. They collect time series and other types of data to help clients like Disney+, Hulu, and The New York Times optimize viewing experiences. Bitmovin tracks every error, buffering, and occurrence affecting user experience to answer questions like: How long does a video take to start? How long are people watching?



As you can imagine, that's a lot of data: CrateDB ingests around 30,000 requests on average, with peaks of 80,000 requests. As Bitmovin queries and ingests data at the same time, queries can easily span 10+ TB. With CrateDB's support for high-availability and automatic failover, Bitmovin can ensure their data is always available, even during peak loads.

140,000 **14**

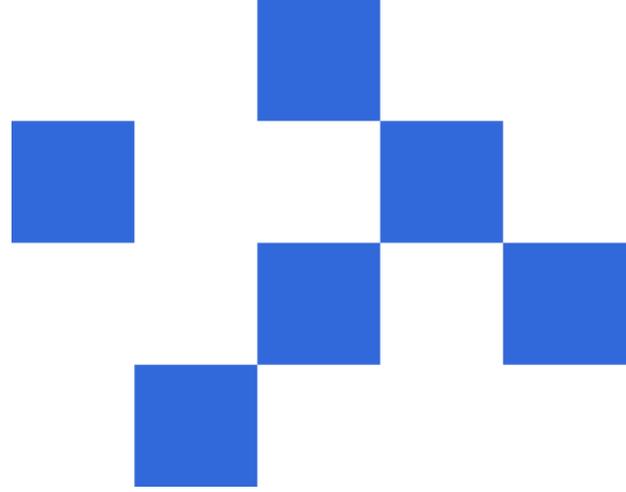
terabytes of storage

Nodes

~60 billion **~2 billion**

events in the largest tables

new events per day



It scales really well. CrateDB scales some extreme loads, and it's very robust and reliable.

*Daniel Hölbling-Inzko
Director of Engineering – Analytics at Bitmovin*

[Click here](#) to check Bitmovin's full story and webinar on how CrateDB provided the backbone of Bitmovin's real-time video analytics platform.

+ CRATE.IO

How Bitmovin improves the streaming experience with real-time analytics

 **Georg Traar**
Head of Customer Engineering
+ CRATE.IO

 **Daniel Hölbling-Inzko**
Director of Engineering Analytics
» BITMOVIN



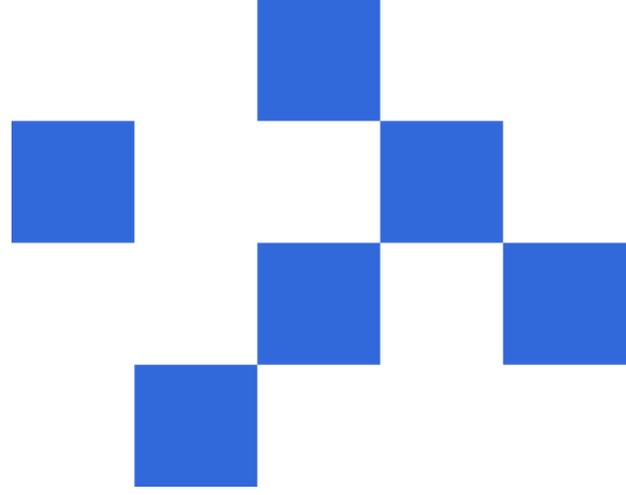
The challenge

Predicting potential issues with critical equipment was a largely unsolved problem in the mining industry but SPGo!'s new solution, powered by CrateDB, has changed all that. Safety inspections in mines tend to be complex, manual, and tiring.

Usually, an inspector walks along the leash with thermal guns measuring and recording the temperature of the poles via photographs and paper. As mineral conveyor belts range from 1 to 15 kilometers, this is understandably problematic.

The solution

SPGo! uses more than 40,000 sensors to monitor mining material conveyor belts in real-time, producing 760 million records per day that are stored in the CrateDB database.



The IoT solution uses predictive maintenance to prevent unscheduled stops which would otherwise interrupt mining operations. Sensors identify sources of potential failures and proactively trigger them during scheduled plant shutdowns, increasing the performance and service life of the machinery and improving the return on investment of costly mining equipment.

30,000

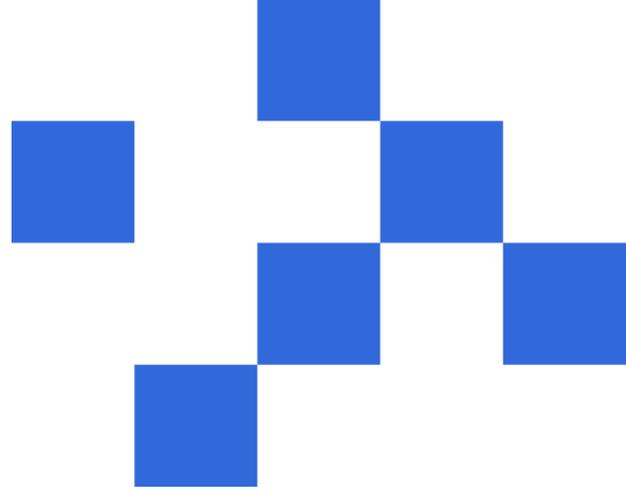
sensors per mine

760 million

records per day

80%

increase in data-driven
predictive maintenance



With CrateDB, we can continue designing products that add value to our customers. We will continue to rely on CrateDB when we need a database that offers great scalability, reliability and speed.

*Nixon Monge Calle
Head of IT Development and Projects*

[Click here](#) to check SPGo!'s full story and webinar on how SPGo!'s Digital Inspection of Conveyor Belts Revolutionizes the Market with CrateDB

January 26, 2023

CRATE

Hernán Cianfagna

IloT technology: How SPGo! builds app for monitoring and predictive maintenance

Hernán Lionel Cianfagna, Senior Solutions Engineer
Michael Mella Navarro, Services and Projects Manager
Nixon David Monge Calle, Development and IT Projects Manager



How can your business take advantage of real-time monitoring?

Real-time monitoring is a valuable asset for businesses of almost all industries and sizes. However, implementing real-time monitoring requires a robust infrastructure that can collect, process, analyze, and report on data in real-time, while also being scalable and secure.

That's where CrateDB comes in: a powerful, scalable database with built-in fault tolerance designed to handle the unique demands of real-time monitoring.

Interested in learning more? Get in touch today for a demo or free consultation. Our experienced team of data engineers will be happy to advise you.

**Get started on your
CrateDB Cloud
FREE trial**

 **CrateDB**

 **CrateDBCloud**