



How to get the most out of your Time Series Data

White Paper





Table of contents

1. How to get the most out of your time series data
2. How can businesses use time series data?
3. What kind of infrastructure does time series data need?
4. Why CrateDB is the optimal database for managing time series data
5. Powering time series data with CrateDB:
Real-world examples
 - 5.1. Bitmovin
 - 5.2. Gantner Instruments
6. How can your business take advantage of time series data?

How to get the most out of
your time series data




Time is **one of the most essential units of measurement**, so unsurprisingly, time series data is of vital importance to almost all types of companies.

Time series data, **also known as time-stamped data**, is a series of data points tracked across time. If you want to measure how much rainfall there is in Brazil each day, the amount of sales orders Nike generates per month, or even the number of heartbeats per minute, you're dealing with time series data.

Time series data is **highly valuable**, but it can be massive and complex, making it a challenge to analyze and manage. To use time series data effectively, you need the right infrastructure to support it. In this whitepaper, we will explore how you can get the most out of your time series data.

How can businesses use
time series data?




Businesses of all sizes and industries use time series data and real-time analytics to track, analyze and make informed decisions based on trends and patterns in their data. This helps organizations improve their operations, minimize risks, optimize processes, and be more profitable.

Here are some examples:

Healthcare

Healthcare related organizations such as hospitals use time series data to monitor patients' vital signs and track their health progress over time. This information is crucial in making informed decisions about treatment plans and identifying potential health issues before they become serious.




Retail

Retail companies track sales and customer behavior, such as foot traffic in their stores and online shopping habits. Time series data helps guide inventory management, demand forecasting, and marketing strategies.

Manufacturing

Manufacturers use time series data to monitor production processes and equipment performance. This helps them identify bottlenecks and inefficiencies in the production line and make improvements to increase efficiency and reduce waste.




Finance

Financial institutions track market trends and analyze financial data, such as stock prices and interest rates. Time series data helps to inform investment decisions and manage risk.

Energy

Energy companies use time series data to monitor the production and consumption of energy in real time. This enables them to optimize their energy grids, reduce waste, and save costs.

What kind of infrastructure
does time series data
need?




Time series data is collected in **massive volumes in real-time**. Real-time analysis and quick retrieval of data are critical, therefore the infrastructure supporting it must be able 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 of time series data, leading to slow and unreliable performance.

That's where a **purpose-built time series database is a real business advantage**. Designed for scale and high-speed retrieval, a time series database like CrateDB delivers speed, reliability, and the power to process vast amounts of time-stamped data in real time.

Why CrateDB is the optimal
database for managing
time series data




CrateDB is a distributed SQL database that makes it simple to store and analyze large amounts of data in real-time. Built in 2012 by a team of engineers, CrateDB helps businesses all over the world to get the most out of their time series data.

Here are some [real-life stats](#) from companies using CrateDB:




- Data flowing from 100+ factories and 1000+ production lines
- 1,000 different sensor message data structures
- 900,000 inserts per day
- 100+ terabytes of retained data
- Real-time dashboards executing ~ 1 million queries per day



CrateDB is [highly scalable](#), so it can handle time series data workloads of any size with ease because of the way it is architecturally designed to store and retrieve data.

As an ANSI SQL DBMS, CrateDB is easy to learn, easy to integrate, and does not lock you into proprietary data access interfaces.

On the next page, find [eight more reasons](#) people choose CrateDB to manage their time series data.



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 optimized for time series data

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

Advanced analytics for time series data

Partitioning, parallel processing, and in-memory columnar indexes enable real-time, complex analytics on time series data, 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, an 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 time series data
with CrateDB:
Real-world examples




Let's dive into two examples of how companies are using CrateDB to manage their time series data.

Bitmovin

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.

Gantner Instruments

Gantner Instruments is a global leader in the development of high-precision measurement and control systems. They provide innovative measurement systems and testing equipment for electrical, mechanical, and thermal measurement.




Bitmovin

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.


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.




“

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

”




Gantner Instruments

With the growth of renewable energy sources, grid operators are facing the challenge of maintaining stable grid frequency, [making constant data collection](#) essential.

Gantner partnered with the University of Cyprus to find better ways of managing electricity consumption and power generation in local grids with high proportions of renewables.


To [analyze the large amounts](#) of data generated by this project in real-time and improve processes through machine learning, Gantner turned to CrateDB, a multi-model database designed for easy data analysis at scale.



With CrateDB, Gantner offers [real-time control and measurement of smart grid services](#) through its Advanced System Monitoring and Analytics solution.

With [CrateDB's ability to handle large amounts of data](#) and support real-time analysis, Gantner is able to provide its customers with advanced analytics capabilities and AI-driven analytics.

Gantner's next-generation solution, powered by CrateDB, enables energy providers to improve system performance, decrease the Levelized Cost of Electricity (LCoE), manage smart grid controls, and ultimately, save costs.




“

Dealing with high-volume and high-frequency sensor data, CrateDB is the only database that gives us the speed, scalability, and ease of use that our teams, customers, and applications require.

Jürgen Sutterlütli,
VP, Energy Segment and Marketing

”

How can your business take advantage of time series data?



Time series data is a valuable asset for businesses across industries. However, its complexity can make it challenging to harness its full potential.

That's where CrateDB comes in: a powerful, scalable, and high-availability database designed to handle the unique demands of time series data.

Interested in learning more? [Get in touch today for a demo or free consultation](#). Our experienced data engineers will be happy to advise you.

Watch now the webinar recording: "Time-series data: From raw data to fast analysis in only three steps."

 **CrateDB**

 **CrateDBCloud**