

ONE PAGE SUMMARY

IT & Software > NoSQL > Time Series Databases

Introduction to Time Series Databases

Unlocking the Power of Time Series Data Management

DEFINITION

Time series databases are designed to handle time-stamped data efficiently. They excel in storing, retrieving, and analyzing data points collected over time.

KEY CONCEPTS

Time-stamped Data

Data entries are recorded with a specific time, allowing for chronological analysis.

High Write Throughput

These databases can handle a large volume of writes per second, essential for real-time data.

Data Retention Policies

Policies that determine how long to keep data, balancing storage costs and data relevance.

Aggregation Functions

Functions that summarize data over time intervals, like averages or sums, to simplify analysis.

Downsampling

The process of reducing the resolution of time series data to save space and improve performance.

EXAMPLES

- Monitoring server performance metrics over time.
- Tracking stock prices at minute intervals.
- Recording temperature readings from sensors every hour.

MEMORY TIPS

- ★ Think of time series as a 'timeline' of data.
- ★ Remember 'write high, store smart' for performance.
- ★ Use 'agg' for aggregation to recall summary functions.

COMMON MISTAKES

- ✗ Confusing time series with regular databases.
- ✗ Neglecting data retention policies leading to storage issues.
- ✗ Overlooking the importance of downsampling for analysis.

QUICK RECAP

Time series databases specialize in managing time-stamped data for efficient analysis. They support high write throughput and include features like aggregation and downsampling to optimize data usage.