Enterprise Historian

Industrial Database Technology

PcVue® Solutions

canary

Mastering interfaces between

#Humans

#Systems

#Connected things

Enterprise Historian

The Time Series Database for Industrial Automation

An industrial data platform designed to help you collect, store, and analyze all your organization's process data.

The Database Technology

The Enterprise Historian is a NoSQL time series database. That simply means it is built to specifically store time series data and does not rely on an SQL database to do so.

Every tag you choose to store in the Enterprise Historian can contain the following data:

- Tag name
- 100+ meta data properties (engineering units, descriptions, limits, and more)
- Timestamps
- Values (booleans, floats, integers, strings, etc)
- Quality score

Imagine the amount of data that can be produced if you have 5,000 tags sending values every second. What if it was 50,000 tags, or even 500,000?

That's exactly why Enterprise Historian chooses not to use SQL, its just too much data for a relational database. The Enterprise Historian was specifically built and optimized for the writing and reading of this specialty time series data.

Incredibly Fast Performance

You can write more than 1.5 million updates per second to the Enterprise Historian in a continuous 24 hours a day, 7 days a week operation. That's a lot of data. Best of all, the database is structured so no matter how many years of data you store, or how many tags you are collecting, you will always maintain that 1.5 million write-per-second performance!

> When needed, the Historian can maintain a continuous read speed of more than **2.5 million reads-per-second.**

The Enterprise Historian gives you better performance, requires no database management, and uses loss-less compression to ensure your data values are never compromised.

BENEFITS

- ✓ No database management required
- ✓ Provides loss-less compression
- ✓ Performance unaffected by database size
- ✓ Highly compressed storage
- ✓ No additional database licensing costs
- ✓ Accept SQL queries

Loss-Less Data Compression

Enterprise Historian's loss-less compression algorithm ensures your data is never compromised. Each day, all your historic records are validated, compressed, and closed for writing. Your original raw data format is forever stored securely and with the smallest storage footprint possible.

Canary achieves an industry leading compression ratio of better than 3:1, saving you more than 66% in storage.

Easy Database Structure

When you deploy Enterprise Historian, you organize your tags into DataSets. A DataSet is a collection of sensors, or tags, that you choose to group together. Since you license the Historian only by tag count, you can create as many DataSets as you need.

Within each DataSet, the Enterprise Historian writes to a Historical DataBase file, or HDB file for short. The HDB contains all the tag names and records the timestamp for every value change as well as its quality score. You can also associate properties to each tag allowing you to store descriptions, engineering units, limits, and more. Typically, a new HDB file is created daily.
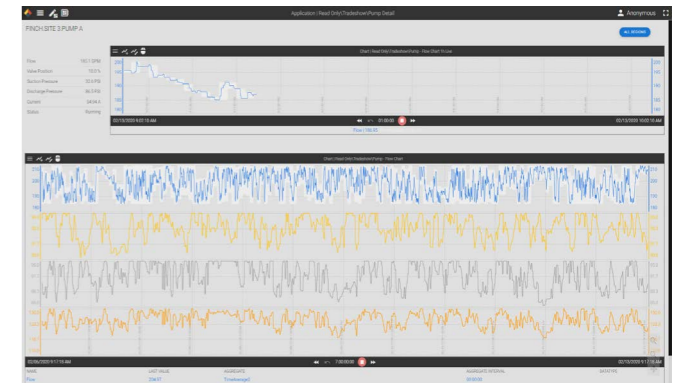
Unlimited Scalability

We know how important it is that your technology scales with your company. That's why both the the Historian's technology and Enterprise Historian's business model are designed for scalability.

A single Enterprise Historian server can scale from only 100 tags to more than 2 million without requiring any additional software installations. Simply adjust the tag license to reflect the number of tags you need to store or even opt for an unlimited tag licensing option for ultimate peace-of-mind.

Architecture, Redundancy, & Disaster Recovery

You can install Enterprise Historians at local sites as well as at corporate locations. Link multiple Historians to automatically move data from the site level to the corporate level in real time or on a schedule. You can also build in redundancy for high availability solutions. Every Data Collector can push data automatically to multiple Historian instances. Additionally, the Enterprise Historian Mirror Service allows you to schedule DataSet snapshots on an hourly, daily, weekly, or monthly schedule allowing for data duplication to offsite Historians.
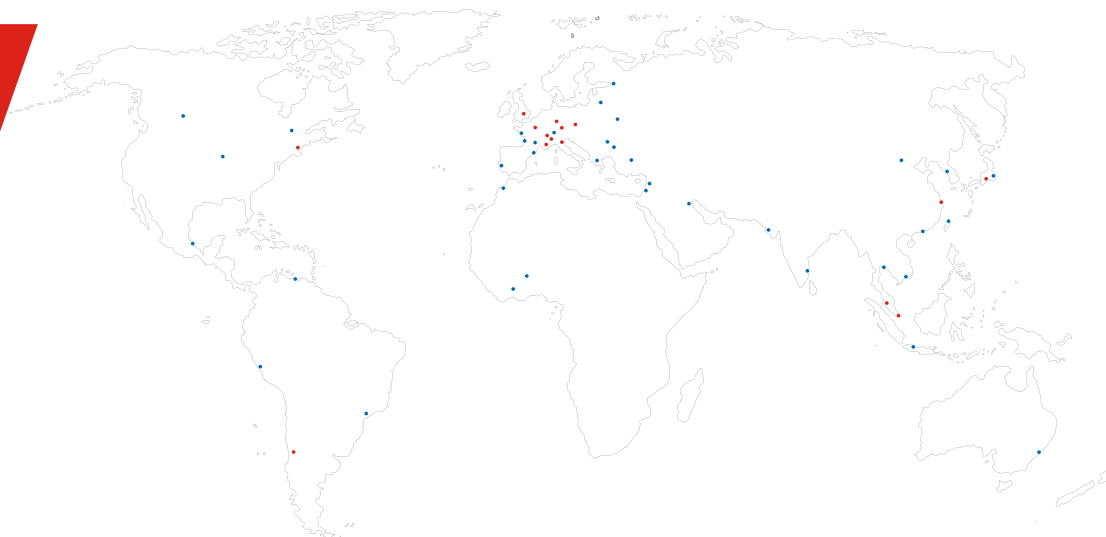
Enterprise Historian™ is developed by Canary Labs

Software platform for IoT, SCADA, BMS & real-time data analytics

ARC Informatique

Headquarters and Paris office
2 avenue de la Cristallerie
92310 Sèvres - France

tel + 33 1 41 14 36 00
fax + 33 1 46 23 86 02
hotline +33 1 41 14 36 25

arcnews@arcinfo.com
www.pcvuesolutions.com

ISO 9001 and ISO 14001 certified

ISO 9001
ISO 14001

Microsoft ®
CERTIFIED
Partner