



Canary System Whitepaper

2024 GUIDE

The Canary System

The Time-Series Data Management Solution Built For Industrial Automation.

Most companies struggle to maximize the value of their process data. At Canary, we believe utilizing process data shouldn't require overtime. Our platform is designed to make using your process data easy. Companies need to store and analyze process data so they can operate more effectively, but they continually encounter the same three problems.

- They produce too much data to manage
- The necessary software is expensive
- Analytic tools are overly complicated

These three roadblocks keep companies from moving forward and leave both engineers and executives frustrated. Canary helps companies defeat these problems. The Canary System is affordable, provides an easy way to store large amounts of process data, and provides simplified analytic tools that deliver powerful results.

Canary Terminology

Canary System - An industrial data platform designed to help you collect, store, and analyze all your organization's process data. Canary has developed multiple software tools that all work together with your existing SCADA solution to provide a complete solution at an affordable price.

Sender - A service that collects, compresses, and encrypts data locally before sending it to the Canary Receiver.

Receiver - A service that moves the data received from the Sender to the Historian for permanent storage. Receiver can communicate with multiple Senders and allows for logging redundancy.

Store and Forward - The ability of the Sender to buffer data locally when the Receiver or Historian is unavailable, preventing data loss.

Datasets - A group of logically associated tags stored within the Canary Historian.

Tag - A single stream of time series data, also referred to as points, channels, or items.

Historical Database (HDB) File - Within each Dataset, the Historian writes tag properties and TVQs to an HDB file. HDB files are segmented by a time-period, typically daily.

TVQ - For each tag value change, Canary records a Timestamp, a Value, and Quality. Other historical databases and SCADA platforms may refer to these as VTQs, VQTs, or Value-Timestamp Pairs.

Step 1: Collect & Store Your Data

Data Collection

Canary collects data local to the source using industry standard formats like OPC UA/DA, MQTT Sparkplug B, SQL, CSV, or through manual data entry. Canary also can collect data directly from SCADA systems through a .NET or Web API. Canary Collectors are used to create logging sessions. Each logging session is configured based on the data source and allows for tag naming, dead-banding, data transformation, and additional configuration.

Once data logging is configured, Canary's Sender Service encrypts, compresses, and caches the data locally. Sender will then push the data packet across the network to the storage portion of the Canary system. You can use multiple logging configurations and Sender instances across your company without any need for additional licensing. In addition, each Sender can push data to additional Canary storage systems creating redundancy and fast disaster-recovery.

To move data for long-term storage Canary uses a Store and Forward Service comprised of the Sender and Canary's Receiver Service. The Receiver Service accepts data from multiple Senders and moves that data into the Canary Historian database. If Sender and Receiver lose contact, or if Historian is offline, Sender will buffer all data locally. Once the Historian's availability returns, all buffered data backfills automatically without any data loss.

Data Storage

Canary has designed its solution to eliminate the need of a database administrator to manage the Historian. Once the system has been originally configured, it needs no ongoing database management. To help achieve this, Canary uses a simple but elegant method to segment the data in two ways.

First is the creation of Datasets, or a group of tags with logically associated data. Second is the Historical Database (HDB) file, segmented by a time period. These two organizational structures allow for database expansion without affecting performance or requiring any management of size or duration.

Generally, a new HDB file is created daily. Once closed, each HDB goes through loss-less compression to minimize Historian's storage requirements. The Canary System never drops, interpolates, down samples, or averages values. Canary's industry leading loss-less compression algorithm preserves the original data values forever.

Each HDB file contains tag names, tag properties, and TVQ units comprised of a timestamp, a value, and a quality score. Values can include Booleans, Integers, Floats, Doubles, or Strings. Canary writes all three TVQ components together ensuring the value is always paired with the correct timestamp and quality. Other databases may store the quality as a separate tag. Canary stores quality with both the timestamp and value to ensure it is never lost or out-of-sync. As an added benefit, additional tag licensing for quality scores is avoided.

The data storage structures of the Canary Historian are designed to maximize read speeds rather than write speeds. Reads are prioritized because data is only written once but read many times. To provision read performance beyond 4,000,000 updates per second, HDB files are organized so TVQs are physically stored with their associated tag, rather than the traditional method of storing TVQs in a time block. As a result, queries made to the Historian on a tag-by-tag basis are optimized for both high tag count and high-speed performance. A single historian can be scaled to store values for two million individual tags, and additional historian servers can be clustered to provide a highly performant and scalable solution capable of handling tens of millions of tags. The Historian performs continuous write speeds at 2,500,000

updates per second and achieves updates rates as fast as ten milliseconds.

To read data from the Historian, Canary uses the Views Service. Views connects to the Historian and allows for client interaction via Windows Communication Foundation (WCF). Access to the Historian, Datasets, or even individual tags can be limited with Views security parameters. Using Windows Active Directory, individuals or user groups can have inherited or explicit permissions for data access.

The Historian can be installed on either a physical machine or a virtual machine and can be hosted in a cloud environment with multi-tenant capability.

Since each Canary Historian can scale to two million tags and handle a variety of data speeds, system requirements can vary. Factors like scan class, change rate, data type, client activity and additional variables can affect system resources. Below are hardware recommendations for a typical enterprise solution consisting of multiple sites each logging data via OPC and Mirroring data to a corporate Canary Historian. If faster update rates are required, increasing storage capacity is recommended.

Collector and Sender Machine

- Dual Core 2.0 GHz Processor
- 8 GB of RAM
- Windows 10 64 bit or greater
- .NET 4.7.1 or greater
- 250 GB HDD @7200 RPM

Site Historian Server

- Quad Core 2.4 GHz Processor
- 16 GB of RAM
- Windows 10 64 bit or greater
- .NET 4.7.1 or greater
- 500 GB HDD @7200 RPM

Corporate Historian Server

- Six Core 2.4 GHz Processor
- 32 GB of RAM
- Windows 10 64 bit or greater
- .NET 4.7.1 or greater
- 4 TB HDD @7200 RPM

Storage requirements can greatly vary. Since the Canary Historian writes changes based on exception, a tag's change rate affects storage. Equally important are the data types being recorded. Finally, the number of years a company desires to store data will adjust requirements.

Therefore, when estimating data storage requirements, it is important to assume significant variations in the types of data, the scan classes of that data, and the change rate.

Step 2: Apply Context To The Data

Canary features several data contextualization and analytics tools: Virtual Views, Calculation Server, and Events. Each of these are included with a Historian installation at no additional charge.

Virtual Views & Asset Models

Most organizations are trying to solve the same problem, how to add context to large volumes of process data. Canary helps companies solve this problem using the Canary Views Service. Primarily, Views receives requests for data from client tools and then extracts that data from the Canary Historian.

When creating a new Virtual View, admins use regular expressions to construct rules that will reshape tag structure and alias tag names. The new view allows for an easy way to standardize tag naming structures without duplicating data in the Canary Historian or requiring the reprogramming of PLCs or SCADA solutions.

Neither the physical data structure within the Historian database nor the available tag license count is impacted by the creation of Virtual Views. The original tag names are not changed, and data is never duplicated.

Additionally, asset models can be defined by creating regular expression templates. Tags are then grouped into assets and can then be browsed by their asset structure. Multiple asset models can be created, and assets can function as subsets, or children, of other assets.

An additional benefit of Virtual Views is the automatic discovery of new tags and assets as they come online. Several Canary Collectors can automatically recognize and add more tags to the Historian on a scheduled basis. As new tags are discovered, the asset models automatically group them based on template rules. For large enterprise systems this greatly reduces the amount of manual work typically required to manage assets.

Canary Calculation Server

The Canary Calculation Server allows for the creation of calculated tags and stores that data in the Historian as if it is being collected from the field in real-time. These calculated tags can be templated on an asset type, and then scaled across an entire group of assets. By leveraging the asset models, admins can define a single calculation and then apply it to every instance of an asset type within the model. Data can easily be backfilled as well.

Beyond basic calculations, logic and aggregated data values based on time, can also be created. This allows for scheduled calculations that not only summarize hourly, daily, weekly, or monthly data, but also makes it possible to 'roll up' data from various assets into totals, averages, and more.

Events

Users may also create condition-based asset monitoring rules using Canary's Events tool. Designed to find and store unique operational events and related data, the Events Service monitors critical processes as they occur. Events can be used to track startups and shutdowns, phases, batches, operator shifts, down-times, and processes that go beyond acceptable limits. Information captured is specific to the duration of the event and provides calculated metrics and key performance indicators throughout the event's duration. All Canary Events are logged in a SQL database and can be accessed by Canary

analytic tools as well as third-party applications.

Step 3: Maximize Your Operation

Axiom

Engineers, operators, and executives use Axiom to build reports, create trends, or monitor their industrial processes. Since it's built using HTML, Axiom works from smartphones, touch screens, laptops, or any modern web browser.

Axiom is easy to understand and requires less than thirty minutes of training for most users. Axiom offers many different visualization tools including feature rich trend charts and HMI dashboards.

Using the trending tool, users can individually build their own trend charts, displaying many tags on a single screen. Once added, tags can be banded together, have scaling adjusted, and display high and low limit alerts. A list of OPC aggregates can be applied to each tag, allowing for aggregated properties to be displayed based on time interval. All chart data can be exported directly from the trend charts to CSV. Additionally, users can create calculated tags displaying values in real time or historically.

Charts can be saved centrally to the Canary server in public or private folders.

Trend charts are just one of the many dashboard elements that Axiom offers. Full multi-screen HMIs can be created using the built-in drag-and-drop editor. Other elements include tables, gauges, symbol graphics, iframes, panels, and more. And since every item can be scripted using C#, advanced workflows and interactions can be created.

Additionally, Canary's Asset Model is fully incorporated into Axiom dashboards and trend charts allowing for asset comparison and monitoring at great scale with very limited work. By creating a report on a single asset instance, Axiom will self-discover all known asset occurrences, duplicate the report for each, and allow the user to define filters assisting in the search for specific conditions.

Canary Excel Add-in

To streamline workflows, Canary offers the Excel Add-in. Users may access data directly from the Canary Historian with software they are already comfortable using. Use the Excel Add-in to import lists of tag names, search for last known data values, or find specific time periods when tag values were outside of norm.

Additionally, they can export large quantities of raw or processed data values or even access Canary Events, running ad hoc asset analysis. No matter what data is needed from the Canary Historian, it's made available within Microsoft Excel.

Data Connectors

For additional data connectivity to the Canary System, use one of the many available data connectors pre-integrated into the Canary System. The Web & .NET APIs allow for custom connectivity to virtually any other system, and clients that utilize ODBC can use the ODBC Connector to expose the Canary System in an SQL-like manner.

Additionally, Canary can be easily integrated with other software solutions to ensure connectivity, including:

- **Spotfire**
- **Flow Software**
- **SEEQ**
- **PowerBI**
- **Citect SCADA**
- **TopView**
- **ESRI**
- **Maximo**
- **and more**

Canary System Pricing

The Canary System includes all the software needed to collect, store, and analyze your organization's process data, and is available as either a subscription or perpetual license.

Each system contains data collectors for OPC UA and DA servers, MQTT Sparkplug B brokers, SQL databases, CSV files, various SCADA systems,

and both .NET and Web APIs. All collectors are free and can be used as often as needed.

Client tools to read and analyze your data are included with each system. These include Axiom, Canary's tool for building trends, dashboards, and reports, as well as the Canary Excel Add-in, which connects Microsoft Excel to the Canary System.

An unlimited number of concurrent Axiom Clients and Excel Add-ins are included with each historian to maximize the utility of the Canary System for interested users across your organization.

Extract data from the Canary Historian using the .Net and Web APIs that are included with each Canary System. Additionally, Canary has partnered with other software providers to ensure you can easily move data between platforms. These include Spotfire, Flow, Seeq, PowerBI, ESRI, Maximo, TopView, and more.

Once you have customized the Canary System to your specific needs, decide where to host your data. You can install the system on your own servers or use the Canary Cloud. If hosting your own data, choose between a perpetual or subscription model.

To instantly view transparent pricing for a Canary System of any size and configuration, visit:

<https://www.canarylabs.com/pricing>

Try Canary

We make using your process data easy. We also make it easy to try Canary. Evaluate the solution for 90 days at no cost.