Canary System Whitepaper

The Time-Series Data Historian Built For Industrial Automation

Introduction

Most companies struggle to maximize the value of their process data. At Canary, we believe utilizing your operation's data should not be difficult. Our platform is designed to make using your process data easy and affordable. As the competitive landscape in your industry evolves, the need to store and analyze process data becomes essential, but most organizations run into the same three problems:

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

These three roadblocks prevent companies from moving forward and leave both engineers and executives frustrated. The Canary System can help resolve these problems. Our software is affordable, provides an easy solution to store large amounts of data, and provides easy-to-use analytic tools that deliver impactful results.

Canary Terminology

To get a better understanding of the Canary System, we provided you with some of the common terms used in our product.

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

Store and Forward - A service responsible for buffering and transporting data from a Canary Collector to its final destination in the Historian.

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

DataSets - A folder structure that groups tags stored within the Canary Historian.

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 - Stands for Tag Value Change. 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.

The 3 Steps of the Canary System

We found that the best way to learn how our technology works is by illustrating the flow of data in three steps. Let's take a look at those steps in more depth.

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 offers a gRPC and Web API for creating custom data collectors.

Each Collector is paired with a Store and Forward (SaF) service. Once data logging is configured, Canary's SaF encrypts and sends the data to its destination. Each SaF service can push data to multiple destinations to support redundancy and disaster recovery. Multiple logging configurations and SaF instances can be used across your company without any need for additional licensing.

The SaF service is also paired with the Historian and is responsible for receiving any inbound logging sessions. In the event that communication between the Collector and Historian is interrupted, the SaF service can buffer data to local disk until the connection can be re-established, avoiding data loss.

Data Storage

Canary has designed its solution to eliminate the need for a database administrator to manage the Historian. Once the system has been 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, which act as folders within the Historian, organizing tags into logical groups based upon their origin. 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 indefinitely.

Each HDB file contains tag names, tag properties, and TVQ units comprised of a

timestamp, value, and 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 both read and write speeds. 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.

To read data from the Historian, Canary uses the Views Service, which acts as a gateway for any client coming in to consume its data. Whether it be raw or processed, the Views service queries the Historian and returns the requested data back to the client.

Access to data can be controlled by enabling tag security within the Identity service. By default, Canary utilizes Windows Active Directory for authenticating and authorizing users who connect. Individuals or user groups can be granted specific permissions for data access. The Identity service also supports OpenID Connect if wishing to use a different provider than Windows AD, and is beneficial if wanting to support an SSO environment.

The Historian can be installed on either a physical machine or a virtual machine and can be hosted in a cloud environment or on-prem.

Estimating the data storage the Historian requires can be difficult as there are many factors that play into it including the number of tags, types of data, scan classes, and data resolution. The number of years a company desires to store data will also affect the amount of disk space needed. The Historian writes by exception when either the value or quality of a tag changes. This places a significant role in minimizing its disk footprint.

Step 2: Apply Context to the Data

Canary features several data contextualization and analytics tools: Virtual Views, Calculation Server, and Events. Each of these is 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. It is also used to create Virtual Views which can alias tag names and define assets. Clients can then request data from the original Historian View or a Virtual View built on top of the Historian.

When creating a new Virtual View, admins use regular expressions to construct rules that will reshape tag structure and define asset types. 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.

Tags can be grouped into assets and browsed by their structure which inherit a child-parent relationship. An asset type may have multiple child assets of various kinds who in turn have their own child assets. Therein, the structure of a Virtual View may vary depending on the company's needs whether it be a flat model, one or two levels deep, or very hierarchical with multiple levels.

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 Service

The Canary Calculation service 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, logical and aggregated data functions 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 Calculation service. Designed to find and store unique operational events and related data, the Calculation service monitors critical processes as they occur. Events can be used to track startups and shutdowns, phases, batches, operator shifts, downtimes, and processes that go beyond acceptable limits. The information captured is specific to the duration of the event and provides calculated metrics and key performance indicators throughout the event's duration. Canary Events are logged in to a SQL database and can be accessed by Canary analytic tools as well as third-party applications using Canary's Read API or ODBC Connector.

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 intervals. All chart data can be exported directly from the trend charts to CSV. Additionally, users can create calculated tags that display values in real time or historically.

Charts are saved centrally on the Canary server in public, read-only, or a user's private folder.

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 concept of asset modeling 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 an Excel Add-in. Users may access data directly from the Historian with software they are already comfortable using. The add-in can import lists of tag names, search for last known data values, or find specific time periods when tag values were outside of the 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 is 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 and gRPC APIs allow for custom connectivity to 3rd party systems, and clients that utilize SQL can use the ODBC Connector to expose the Canary System in a SQL-like manner.

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 brokers, SQL databases, CSV files, various SCADA systems, and both gRPC 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.

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.

To instantly view transparent pricing for a Canary System of any size and configuration, visit: https://www.canarylabs.com/pricing