Hewlett Packard
Enterprise

Enabling Data Management at the Edge in Industrial IoT

HPE Edgeline IoT Systems & Converged Edge Systems with OSIsoft PI System

Technical white paper

Contents

Executive Summary

One of the most important and challenging tasks facing a company's business and operations teams is how to get the most efficiency and productivity from the product creation and delivery process. Whether a company builds cars, produces food and beverages or explores for and delivers energy products, there is a need optimize the business operations at each step of the delivery process. One of the key ways to accomplish this is by collecting and analyzing data from the operations process. By continually examining data collected from a variety of sources such as the machines used to manufacture products and the staff doing the work, opportunities for new efficiencies can be identified and potential catastrophic problems can be avoided. To achieve the most timely and effective results, this data collection and analysis is often best done in real-time in close proximity to the source of the data, or at the edge. The edge is the location beyond the traditional datacenter where the things generating the data reside. This could be a factory, an oil well, a vehicle or a distribution center. The edge is where the things are and where the data is created.

HPE's new Edgeline IoT Systems and Converged Edge Systems products were designed to be the platforms of choice for the emerging intelligent edge. They provide the data ingest, compute, storage and management capabilities required to perform data collection and analysis at the edge.

OSIsoft provides the PI System software, a data infrastructure that gathers and stores data from a wide variety of operational equipment and systems, providing a realtime and historical view into operations.

The combination of HPE's Edgeline systems and OSIsoft PI System enables operations teams to gain new and timely insights into their processes.

The benefits of this solution are that it will enable better process and machinery monitoring, improve productivity of equipment and employees and help eliminate waste. The combination of these will give a manufacturer better insight into how to best optimize their manufacturing process, decrease operational cost, improve resource management and increase revenue. In order to accomplish these goals data is leveraged from connected devices, people, and machinery in the local environment (e.g. factory, refinery, oil well and distribution center). This information is gathered, aggregated, and processed on-site to generate actionable insight and enable greater control.

This document describes the context, fit and value of HPE hardware and OSIsoft software for managing and processing data at the edge, enabling the deeper insights that customers need to improve their business.

Solution Overview

Customer operations can generate a staggering amount of data from varying devices with multiple data types and formats. To successfully make use of this data, a set of tools must be utilized that can ingest, process, analyze, store and display this data in a timely and productive manner.

OSIsoft PI System is a scalable suite of software to collect, analyze, store, visualize, and share large amounts of high-fidelity, time-series data from multiple sources to people and systems. This data is often scattered among different incompatible systems, formats, and processes. OSIsoft PI System can connect to hundreds of interfaces and collect high-frequency, high-fidelity data from multiple formats, standards or conventions – both time-series and event-based – and translate that data into a uniform structure for combining, comparing, contextualizing and leveraging information.

Combining the HPE Edgeline family of products with OSIsoft PI System allows the collection and analysis of this data to occur at the edge where the data is generated. This reduces the latency between generation and collection of the data, allowing a quicker time-to-action using insight gained from the data analysis. Network bandwidth and cost are also reduced as high-frequency, high-fidelity, and high-volume data is analyzed locally in near real-time, with insightful results and actions shared upstream.

Two scenarios are presented in this solution. The first uses a single HPE Edgeline EL20 running the PI System to collect data, analyze it, store it and forward it to a remote Enterprise PI System. This utilizes different PI interface technologies to easily achieve this connectivity, transmitting data between the data source and the local PI System to access real-time process data from the plant floor. Additional PI interface technologies use standard networking protocols for transferring data between the local PI Server and the Enterprise PI Server. This architecture can be scaled out to multiple EL20 systems, each collecting data using a standard PI interface technologies and sending that data to single receiving PI System for thorough analysis and visualization.

The second scenario uses an Edgeline EL1000 system hosting a single HPE ProLiant m510 running the PI System to collect data and perform visualization using several PI System Tools. Here, PI interface technologies are used to collect data from several dozen files, each containing a year's worth of data collected from various compressors. In addition to collecting data, PI System visualization clients can be used to provide data access, visualization, and analysis to the collected data.

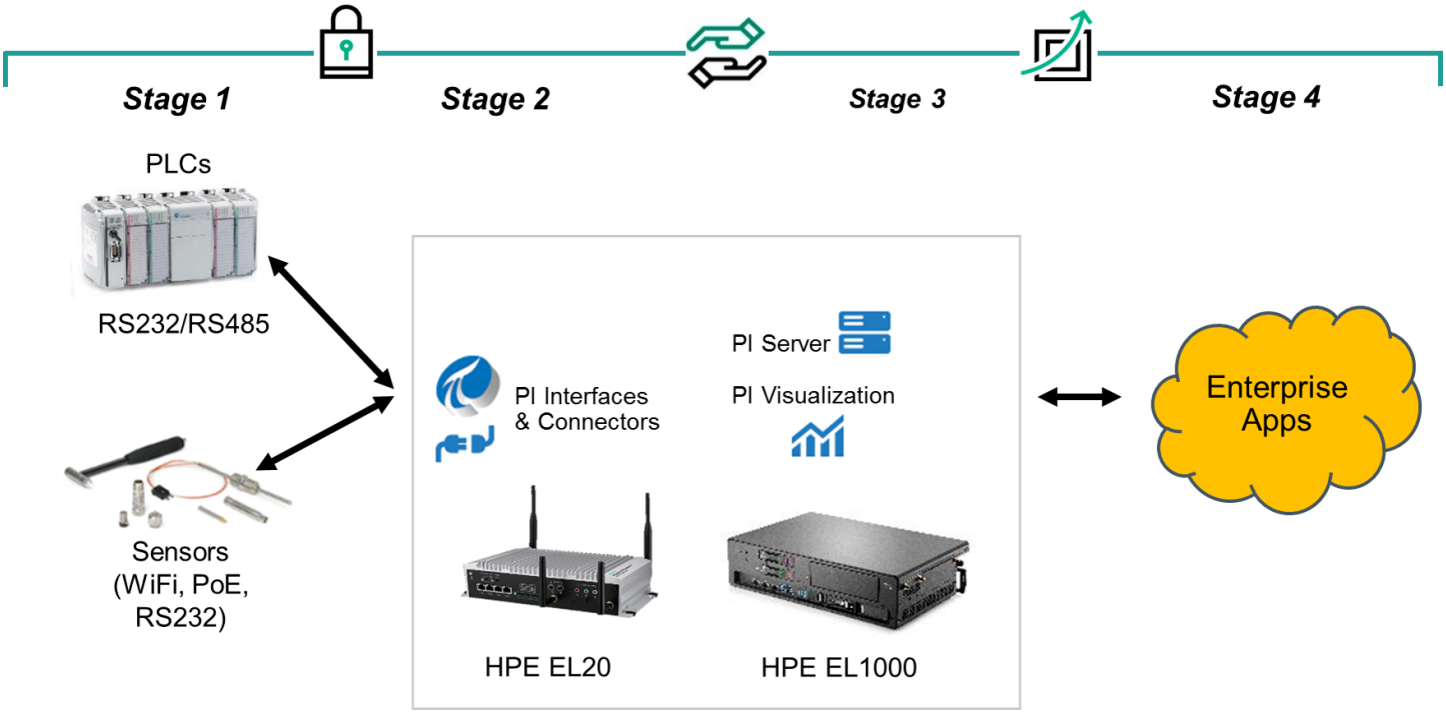

Figure 1. HPE Edgeline and OSIsoft PI System

Solution Hardware Components

This solution utilizes the HPE Edgeline family of products as the foundation for an edge data management solution. The need to quickly analyze and drive business decisions based on real-time data accentuates the need for edge computing. HPE has created unique systems that are purpose-built for converging real-time data acquisition, enterprise-class computing, and remote manageability. HPE Edgeline systems are energy-efficient, ruggedized platforms with a broad range of network connectivity and data acquisition options to accommodate even the most complex industrial application. Chassis type, number of servers, number of CPU cores, memory, and storage can all be tailored to site requirements.

The specific Edgeline products used in this solution are described below.

- **HPE EL20:** The HPE Edgeline EL20 IoT System is a mid-level ruggedized compute solution designed for light data aggregation and acquisition. This Edgeline device delivers a performance solution that comes optimally configured with an Intel® Core™ i5 processor, 8 GB RAM, 64 GB SDD Storage, Wi-Fi connectivity and an expansive I/O selection including four-port Power over Ethernet (PoE) plus 1 x 8 bit digital input/output (DIO).

Figure 2. HPE Edgeline EL20 IoT System

- **HPE EL1000:** This rugged, compact Converged Edge System is designed specifically for harsh environments, providing data center-level capabilities at the edge that delivers immediate insight from IoT data. It combines datacenter class compute cartridges with Intel Xeon x86 processors, performs unique integration of precision data capture and control, and is managed with data center class security and systems management software. In addition it can carry wireless and 3GPP data cards for backhaul traffic. Deep edge compute capabilities will enable businesses to make real-time decisions, adding value to their operational processes that result in better business outcomes.

Figure 3. HPE Edgeline EL1000 Converged Edge System

- **HPE ProLiant m510 server cartridge:** The HPE ProLiant m510 server cartridge is designed to enhance the performance of many general purpose workloads. The ProLiant m510 server cartridge has one Intel® Xeon® D-1548 (8-core) or D-1587 (16-core) with up to 128GB of ECC protected memory, dual 10Gb Ethernet along with up to 2 (1TB NVMe each) M.2 flash storage modules, and up to (1) 240GB SATA M.2 for local OS booting. This is used as the compute engine for the EL1000 system described earlier.

Figure 4. HPE ProLiant m510 server cartridge

Solution Software Components

This solution leverages the OSIsoft PI System software to perform the data collection, analysis, storage, and visualization. The PI System gathers data from a wide variety of manufacturing systems and operational equipment like pumps, compressors, and other industrial assets and then delivers it up in a coherent, unified way so users can optimize their operations, save money, improve safety or develop critical strategic decisions based on analytics. It's a real time view into operations for both real-time and long- term tasks.

For example, the analytics capability of the PI System can flag performance problems with a pump before it breaks. In another use case a user can study the entire history of a manufacturing process to identify the factors most critical for optimizing production and profits.

The ultimate goal is to give people the power to tap into an incredibly important resource—their operational data— that they probably don't use to its full potential.

The specific OSIsoft software products used in this solution are described below.

- **PI Server:** This is the foundation of the PI System, surrounded by technologies that work together to optimize data storage, transformation and delivery throughout the enterprise. This includes sophisticated calculation tools that transform raw data, combine data values from multiple systems into robust analyses, automate calculation of vital day-to-day metrics, and deliver new insights.
- **PI Interface Technologies:** The PI Connectors and PI Interfaces enable the collection of real-time data and includes more than 280 different connectivity options, enabling the PI System to access a diverse set of operational and business data sources including SCADA, DCS, PLC, industry standards, databases, and text files. PI System APIs can also be used to collect data and deliver data to custom applications and share data outside of corporate boundaries using cloud technologies.
- **PI Visualization Tools:** These tools take the complexity out of data access, visualization and analysis. They provide self-service access to the data where and when it is needed.
- **PI Integrators:** These format PI System data for enterprise applications and cloud services to reduce the time and complexity of business analytics. PI Integrators are directly targeted at eliminating the "data janitorial" problem bogging down many organizations.

Specific Use Cases Addressed

Scenario 1: HPE Edgeline EL20

The first solution scenario utilizes a single HPE Edgeline EL20 running the PI System at the edge. From this system, data can be collected, analyzed, stored and forwarded to an Enterprise PI System for further analysis, aggregation and presentation. Multiple EL20 systems could be placed throughout the environment, collecting data directly from dispersed systems, and aggregating data on a single PI System. The EL20 is configured with the following OSIsoft components running on a Windows Server 2012 platform:

- PI Connector for Universal File and Stream Loading (UFL) – provides a file-based and a RESTful interface and a parsing engine to convert data in multiple file and Web friendly formats into the PI System on the EL20
- PI Data Archive – stores time-series data on the EL20
- PI Asset Framework (AF) – stores metadata and organizes this data in the PI System on the EL20 so it's easy to search, share, compare and analyze
- PI to PI Interface – transfers data from the PI Server on the EL20 to an Enterprise PI Server, where multiple EL20-based PI Servers can be aggregated

In this configuration, the PI System on the EL20 can handle up to 15K PI tags, 3KAsset Framework (AF) elements, 30K AF attributes, analytics generating 3K new PI tags, 15K events/second into the PI Server and 15k events/second transferred to an Enteprise PI Server

This example is useful when you need to have PI Systems collecting information from widely distributed devices. For example, if you needed to collect information on a series of geographically distributed oil wells, then placing an EL20 running the PI System at each well might be an excellent choice. The data from these PI Systems could be aggregated on a periodic basis at a more centrally located system such as an HPE Edgeline EL4000 in a regional office for that oilfield.

Scenario 2: HPE ProLiant m510 Server Cartridge

The second solution scenario utilizes a single HPE ProLiant m510 Server Cartridge, running at the edge in the HPE Edgeline EL1000 Converged Edge System chassis. The EL1000 with the m510 cartridge provides a powerful combination of compute, memory and storage, facilitating the real-time collection and analysis of large amounts of data using the PI System.

For this scenario, there are two virtual machines running on the m510 under Microsoft Hyper-V and Windows Server 2012 R2; the PI System software can be installed as needed to meet the desired system architecture.

The server VM is configured with the following OSIsoft components running on a Windows Server 2012 R2 platform with SQL Express 2012:

- UFL Connector – converts file-based data into the PI System
- Data Archive – stores time-series data
- Asset Framework – organizes data in the PI System so it's easy to search and share
- PI Coresight – provides web-based visualization of data in the PI System
- PI Analysis Service – executes calculations defined in PI AF in real-time.

The client VM is configured with the following OSIsoft components running on a Windows 8.1 platform and Microsoft Excel:

- PI Asset Framework client – provides programmatic access to the remote PI System
- ProcessBook – graphics package to create dynamic, interactive graphics from PI System data
- DataLink – an Excel add-in that retrieves PI System data directly in to Excel spreadsheets
- PI Builder – an Excel add-in for publishing tags into the PI System

These two virtual machines could easily be placed on separate physical systems based on the desired architecture. For example, the server VM could be the sole configuration on the m510 and the client VM could be a separate client system on the network.

Data Initialization

The PI System supports the collection of data from a variety of sources and formats. It uses a set of PI Connectors and PI Interfaces to acquire the data from external sources. For some systems and devices, CSV files are a common way of exchanging information. The PI System can be configured to periodically import data from systems that use CSV files and perform analysis. For this example, CSV files are used to acquire data from a set of compressors.

In this scenario, a year's worth of data was collected from several dozen compressors and stored into separate comma-separated value (CSV) files. These files are the data source to be collected and analyzed by the PI System. To collect this data, two things are done.

First, the PI data tags are initialized into the PI Server. Tags are named data elements that store values over time for each desired attribute of the compressor, e.g. flow rate, temperature, etc. There are six tags per compressor which are all defined in an Excel spreadsheet. Using the OSIsoft PI Builder add-in for Excel, the tags are published into the configured PI System directly from the spreadsheet. Recall that the PI Server is running on a separate server system; the client system is configured with the location and credentials of the PI Server on the server system.

After the tags are initialized, the time-series data values can be collected from the CSV files. A PI UFL Interface on the server system is configured to process these data files. The interface defines what the data looks like, how to parse the CSV files, how frequently to scan for new data files, etc. In many environments, using a data file is a common method for continuously gathering data, and in such cases the scan period would be utilized. After all the data values are collected (over 500,000 rows of time-series data per compressor) the data is available in the PI Server for additional analytics.

Data Analytics

Data storage in the PI System is performed using the PI Data Archive. This component provides efficient storage and archiving of time-series data, and high performance interfaces for data retrieval. The PI Asset Framework (AF) builds on the PI Data Archive to integrate and contextualize the time-series data into asset-centric models which simplify data analytics and visualization. These asset-centric elements are available to PI DataLink, PI Coresight, and PI ProcessBook where they can be further analyzed and presented.

PI System Explorer is a graphical tool for managing PI AF elements and attributes. An element template is first created to define the basic structure of the element, i.e. its data attributes such as temperature, flow rate, etc. PI AF provides the flexibility to define an element as a specific system within and industrial asset, the asset itself, a collection of assets, or any other logical or physical grouping of data. Afterwards, individual elements are instantiated based on the template for each physical asset. For our example, four element templates were created to represent four common types of compressors in use, where each element groups all the data together for each compressor type. Element templates allow the creation of element-relative displays where the data from the instances of each template can be easily viewed from a common display. The PI System Explorer additionally allows examination and some analysis of the data. By navigating the element tree, individual attribute values can be displayed. By selecting the Time Series Data or Trend options, all of the archived data can be quickly and easily viewed in either tabular or line chart formats.

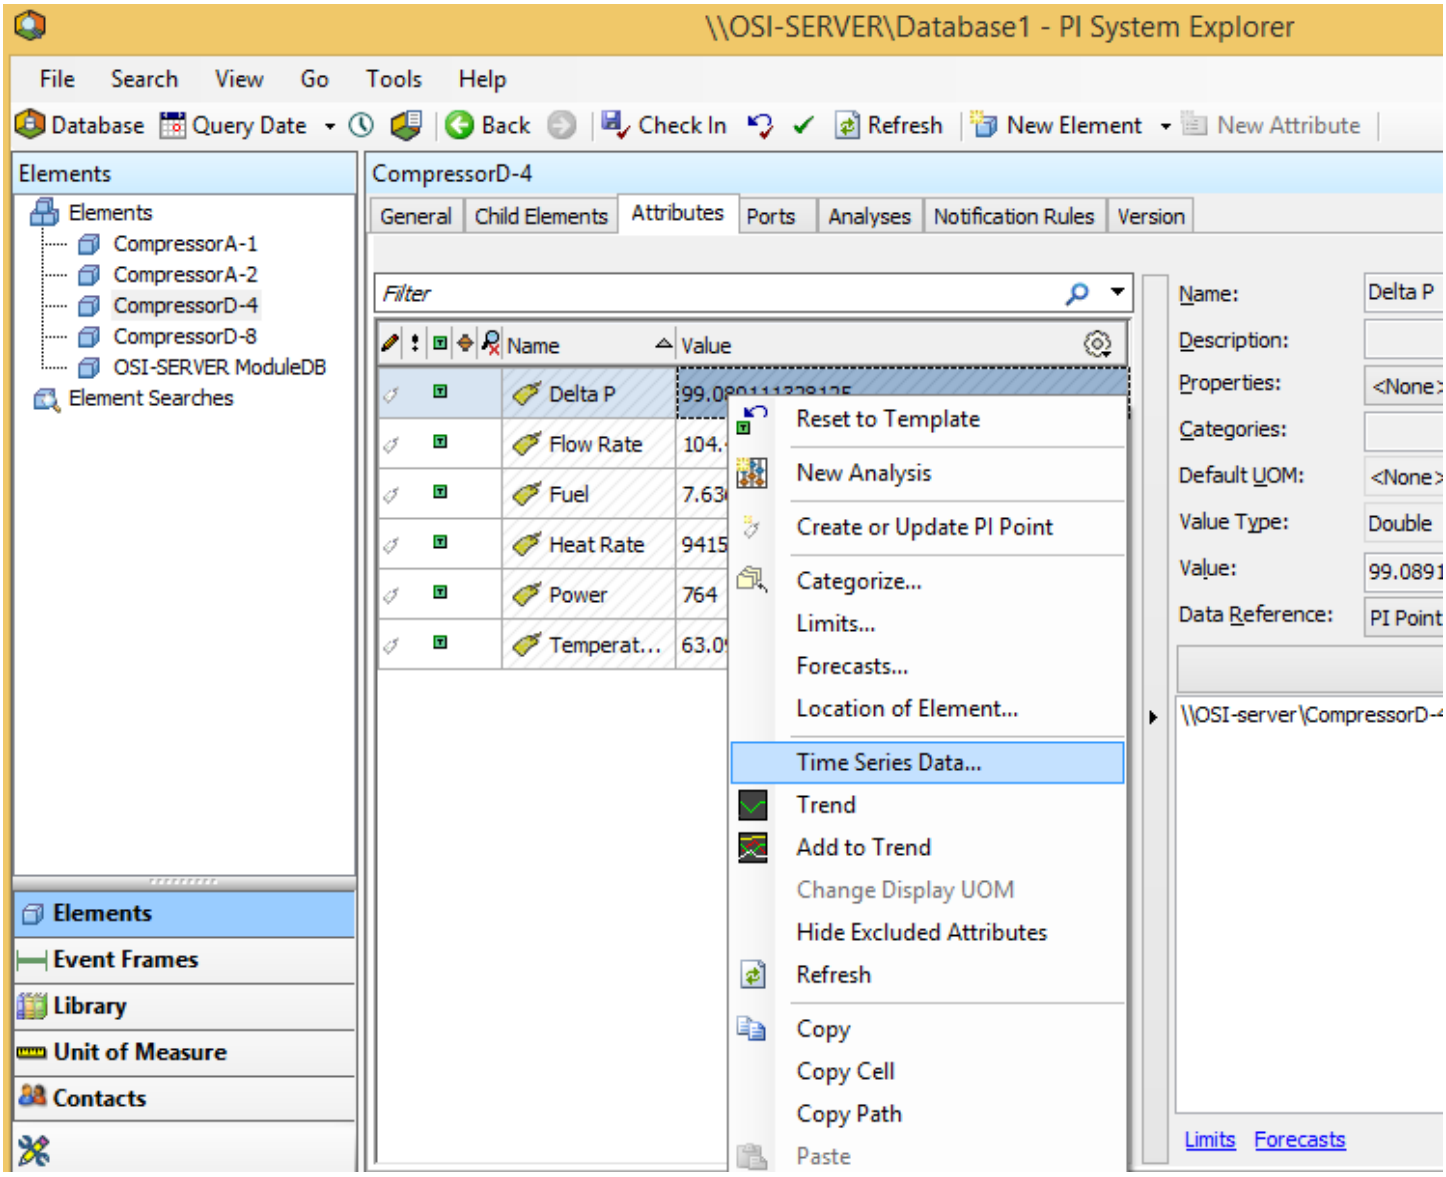

Figure 5. PI System Explorer

PI DataLink is an add-in for Microsoft Excel that allows PI data to be retrieved directly into Excel spreadsheets. Spreadsheets are frequently used to provide periodic reports (monthly, weekly, daily) and PI DataLink makes PI System data easily available to Excel and its analysis and visualization capabilities. The spreadsheet can be designed to pull a fixed set of data values, or it can be flexible to grab values from a relative time frame, e.g. “yesterday” or “last week”, such that simply re-running the spreadsheet daily or weekly produces an updated report using recent data, without needing to manually update each time period’s data. Similarly, using the element relative displays, the data source can be varied to perform reports against different compressors by changing one cell in the spreadsheet to point at a different compressor.

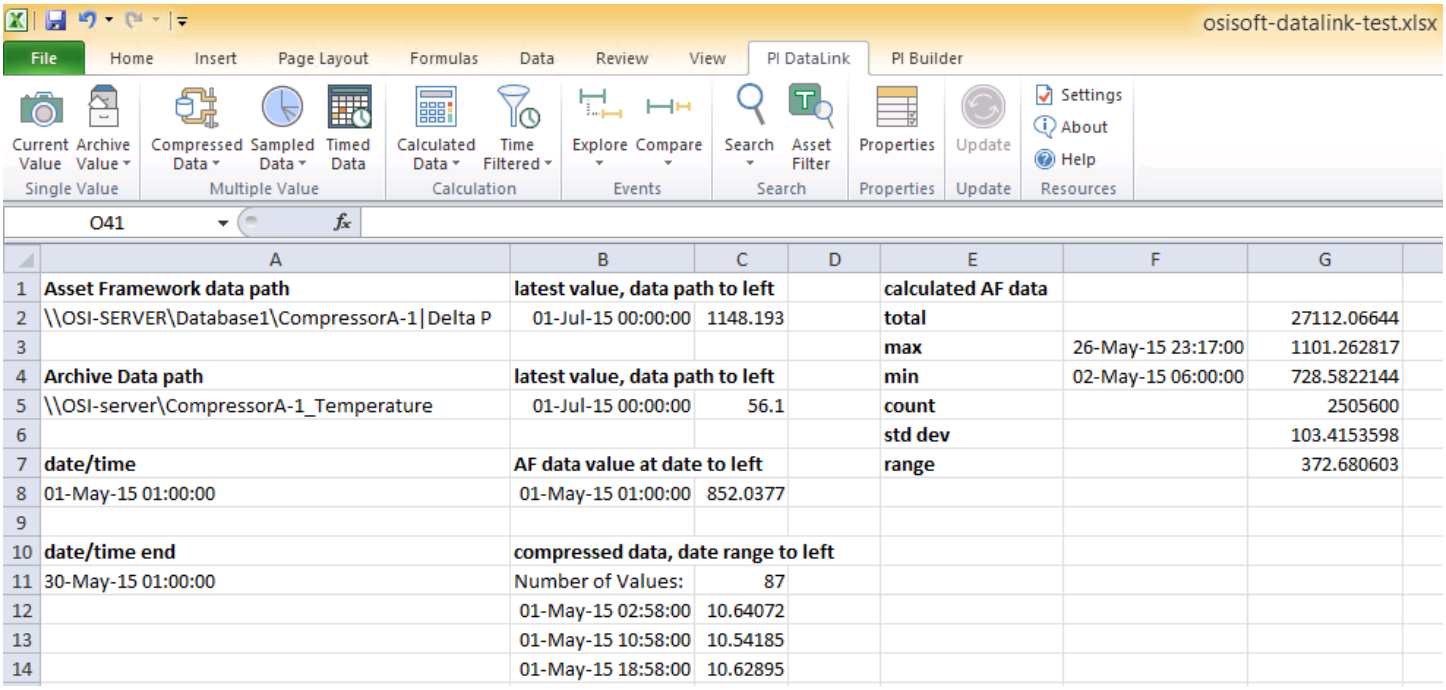

	A	B	C	D	E	F	G
1	**Asset Framework data path**	**latest value, data path to left**			**calculated AF data**		
2	\\OSI-SERVER\Database1\CompressorA-1|Delta P	01-Jul-15 00:00:00	1148.193		**total**		27112.06644
3					**max**	26-May-15 23:17:00	1101.262817
4	**Archive Data path**	**latest value, data path to left**			**min**	02-May-15 06:00:00	728.5822144
5	\\OSI-server\CompressorA-1_Temperature	01-Jul-15 00:00:00	56.1		**count**		2505600
6					**std dev**		103.4153598
7	**date/time**	**AF data value at date to left**			**range**		372.680603
8	01-May-15 01:00:00	01-May-15 01:00:00	852.0377				
9							
10	**date/time end**	**compressed data, date range to left**					
11	30-May-15 01:00:00	Number of Values:	87				
12		01-May-15 02:58:00	10.64072				
13		01-May-15 10:58:00	10.54185				
14		01-May-15 18:58:00	10.62895				

Figure 6. PI DataLink Add-in

PI ProcessBook is a graphics package that integrates PI System data into dynamic, graphical visualizations and trends. PI ProcessBook also incorporates Microsoft Visual Basic for Applications (VBA) to help automate additional activities. PI ProcessBook has a large library containing hundreds of graphical symbols that can be laid out to represent the physical asset and its environment. These symbols are dynamic and can be configured to visually stand out when extra attention is needed as determined by the data analysis. Additionally, ProcessBook takes advantage of the PI Asset Framework to make element-relative data available.

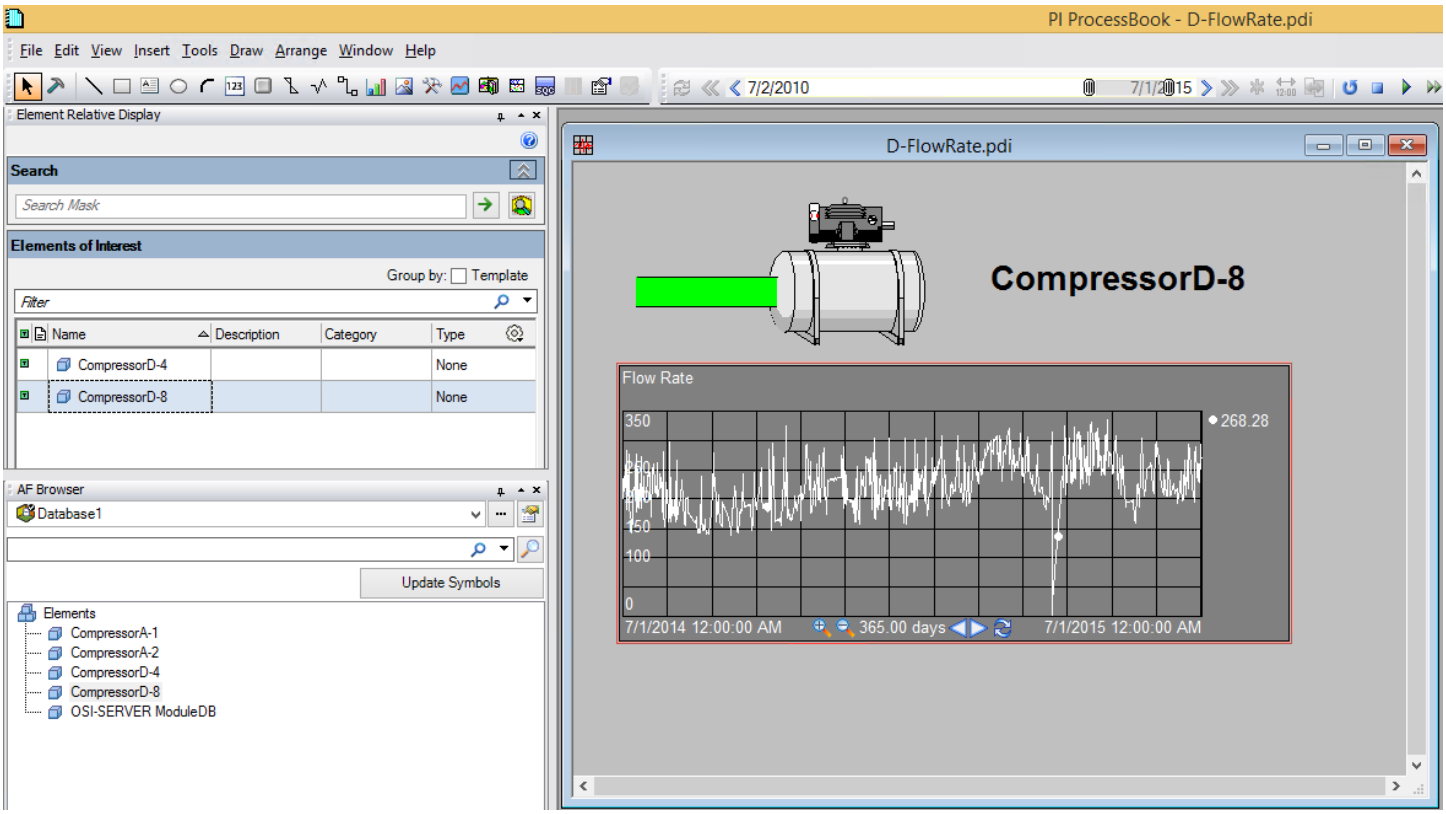

Figure 7. PI ProcessBook Display

In this example, a ProcessBook display (the graphical page containing PI data) is created with a trend graph containing a year's worth of flow rate data, and above the graph is a symbol of a compressor. This page was created using element-relative data such that the compressor source can be dynamically changed to point to another compressor. Note the two compressor entries shown in the "Elements of Interest" section in the figure above. Selection of either of these compressors will cause the page to be updated using that compressor's data.

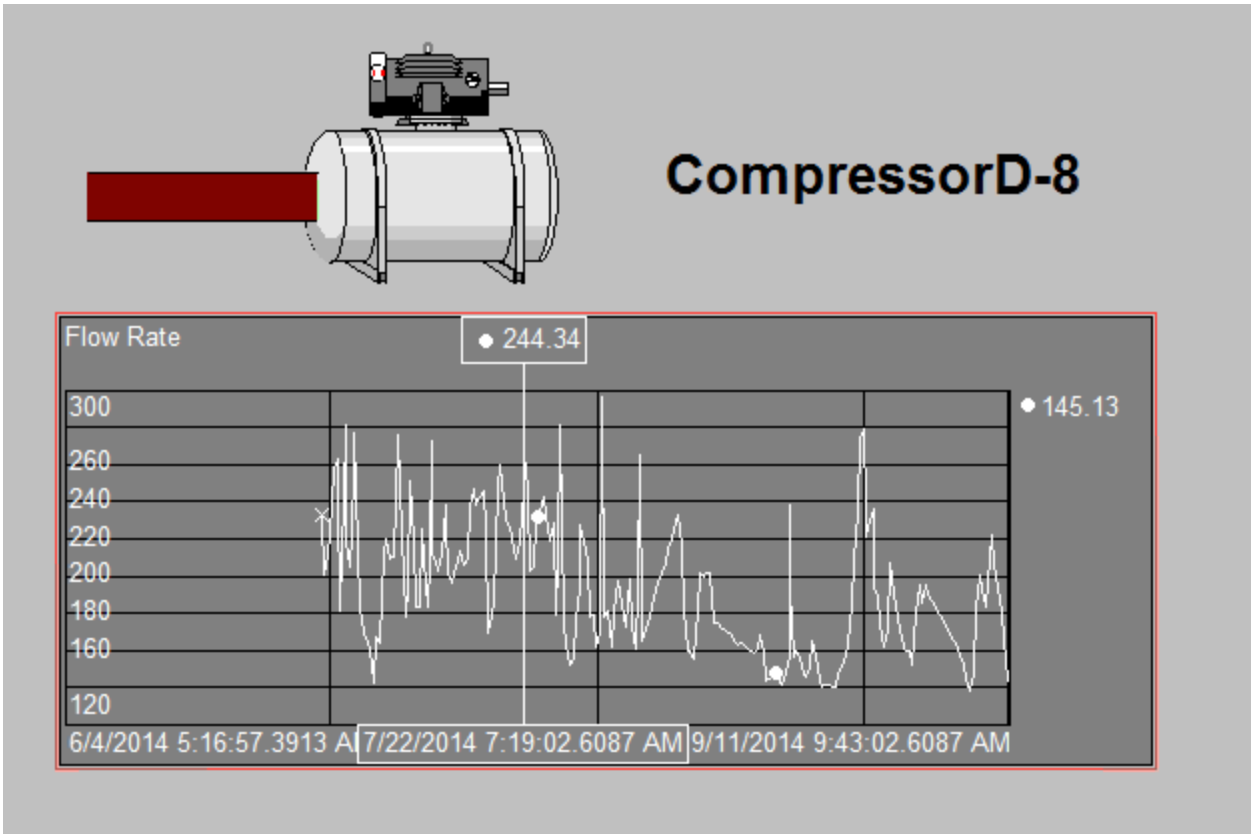

Figure 8. Time Range and Trend Cursor

The “pipe” on the left side of the compressor symbol is a dynamic multi-state symbol that will change colors based on the flow rate to visually indicate if the flow rate is low, high, or within acceptable parameters. If the value is too high, the pipe symbol will also blink. There are several methods for traversing the archived historical data sourcing this page. Using the time range toolbar, the overall range of data can be altered for the entire page. This will also affect the data sourcing the active pipe symbol and trigger the different colors. A low flow rate value is shown in the figure above. To more closely examine data points on the trend graph, a trend cursor can be selected and used to select points on the graph and view specific values. This is also shown in the figure above.

PI Coresight is a web-based visualization tool allowing ad hoc analysis and sharing. As with ProcessBook, custom display pages are created to display and analyze data. Asset Framework data can also be utilized to easily switch data sources between different assets.

Figure 9. Trend Charts

In this example, PI Asset Framework (AF) elements are used so that different compressors can be selected for analysis and viewing. Note the asset selection drop-down control at the top of the page which enables the same display to be used for each compressor of the same type; this demonstrates the use of element relative displays. The contents of the page include the six attributes available for each compressor. This data is included in a table with some statistical analysis performed on the data set. Below the table are a few controls showing only temperature data – an analog gauge and a trend chart. The gauge is configured as a multi-state gauge so that the color will vary based on the temperature value.

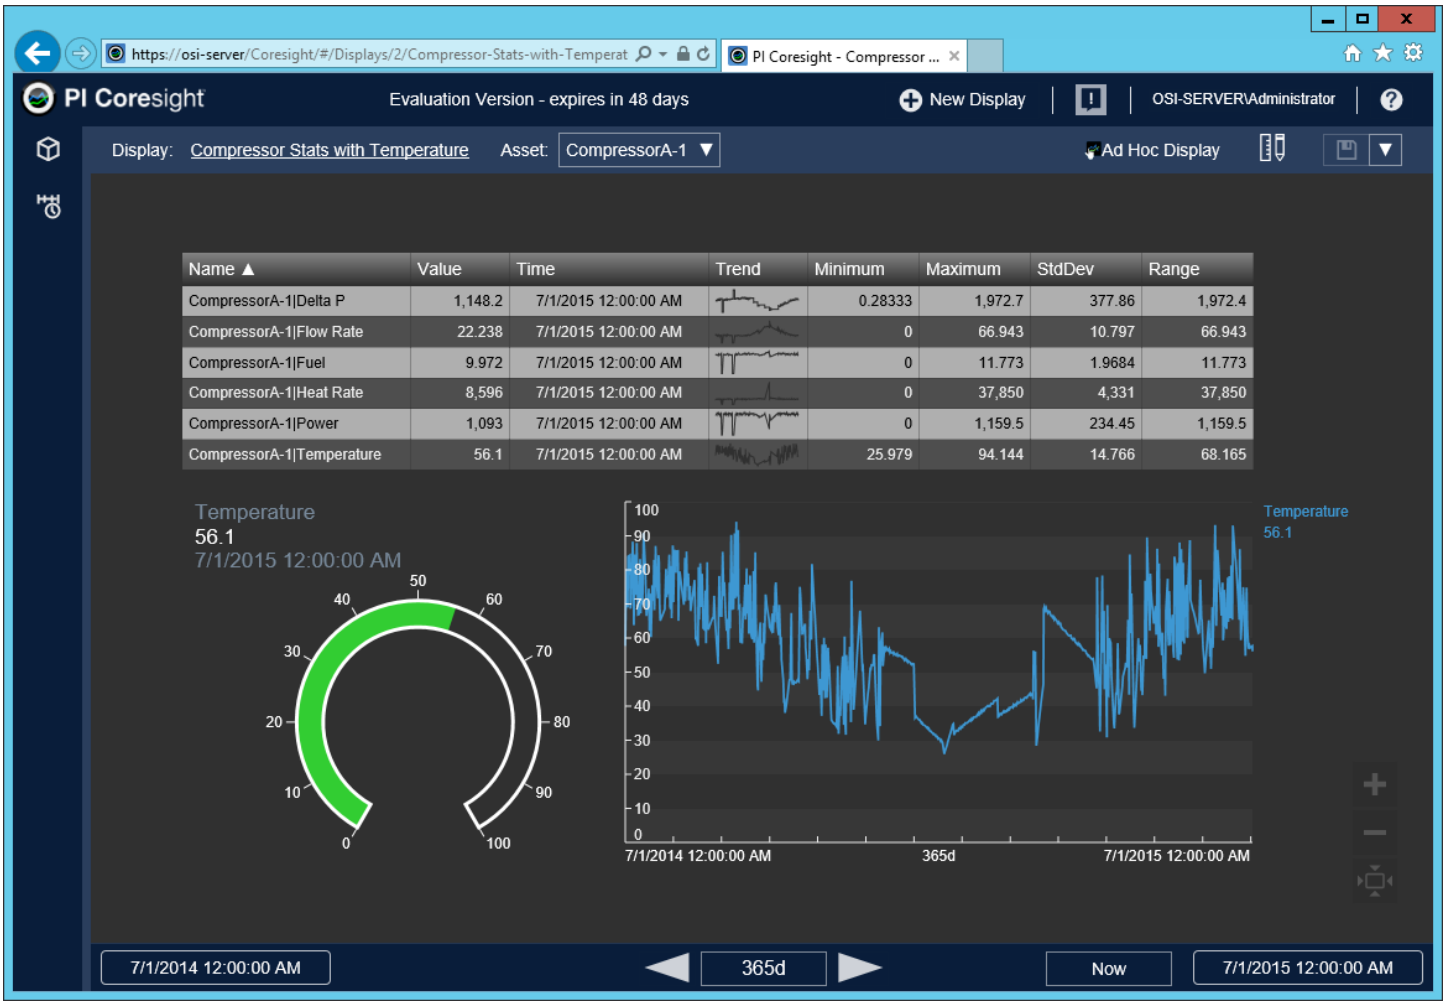

Figure 10. Statistics with Temperature

In this example, Asset Framework elements are used so that different compressors can be selected for analysis and viewing. Note the asset selection drop-down control at the top of the page. The contents of the page include the six attributes available for each compressor. This data is included in a table with some statistical analysis performed on the data set. Below the table are a few controls showing only temperature data – an analog gauge and a trend chart. The gauge is configured as a multi-state gauge so that the color will vary based on the temperature value.

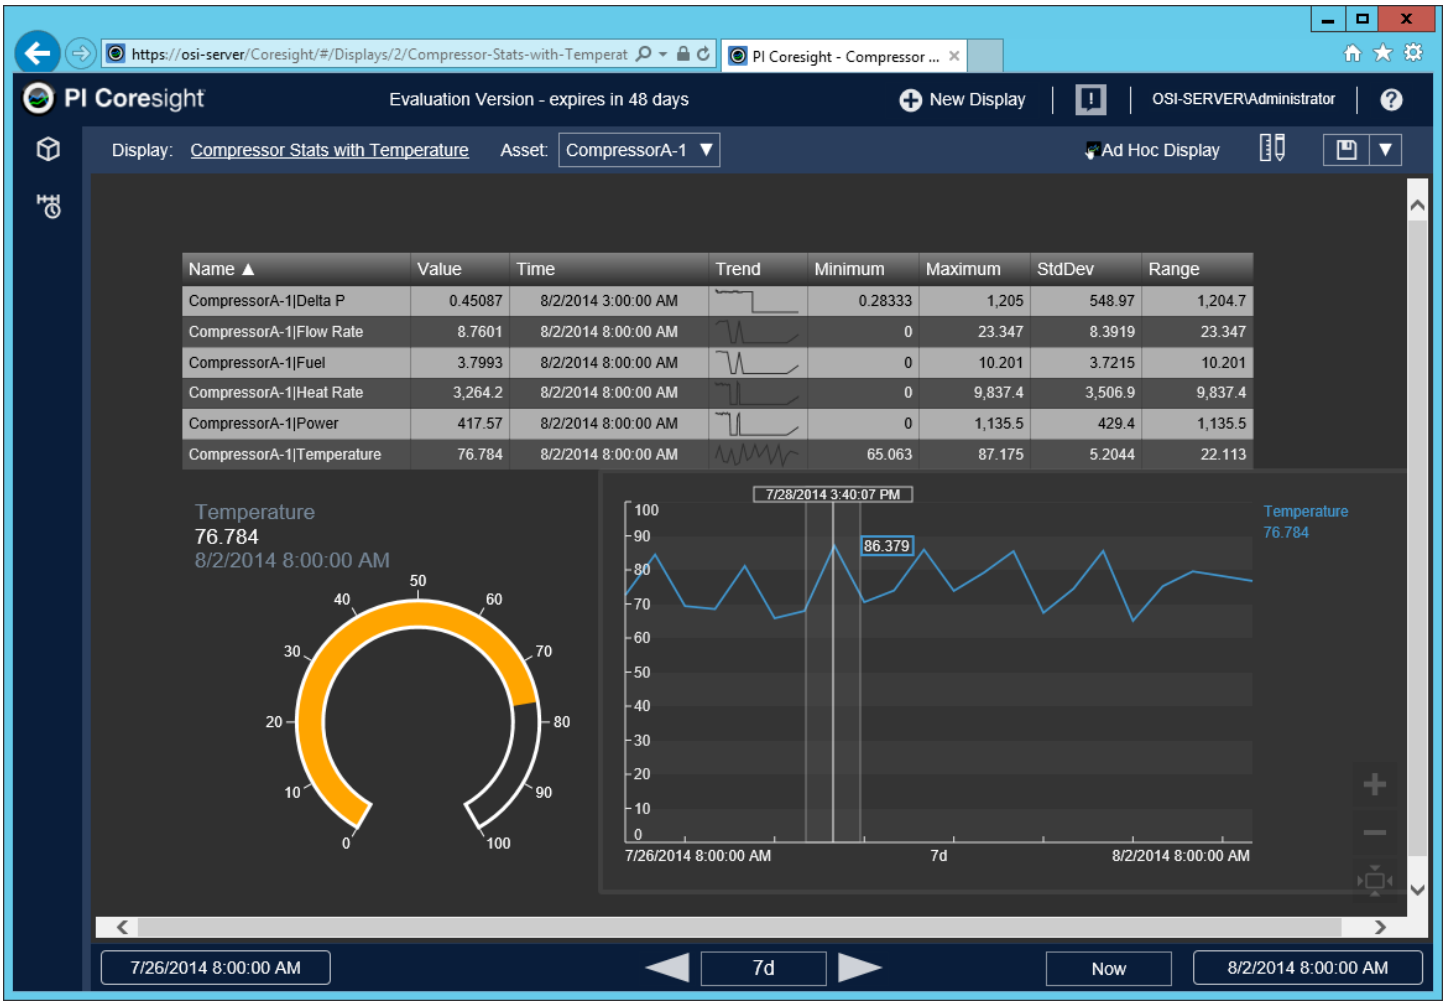

Figure 11. Modified Time Range

The time range controls at the bottom of the page control the data set used for the entire page. In the example above, the range has been reduced down to one week's worth of data. Note that the statistical table has changed since it is now operating against a subset of its previous data. Similarly the analog gauge has now been triggered to indicate the temperature is a little high. Also note the trend cursor in the trend chart that can be used to show details of a particular point in the graph.

Summary

This paper shows how a combination of HPE Edgeline IoT systems and OSIsoft PI System software can be used to collect different data types from multiple sources and perform data analytics and visualization at the edge, improving data management at the edge, resulting in a quicker time-to-action and higher operational intelligence. By performing ongoing analysis of both current and historical data, companies can get the most efficiency and productivity from their product creation and delivery process. This use of HPE's Converged Edge Systems and OSIsoft's PI System enables the deeper insights that customers need to improve their business.

Learn more at

hpe.com/info/edgeline

873729-001, November 2016

Hewlett Packard Enterprise