Supercharging PI System with AI-based PdM

Expanding PI System's Real-time Analytic Capability with Amber

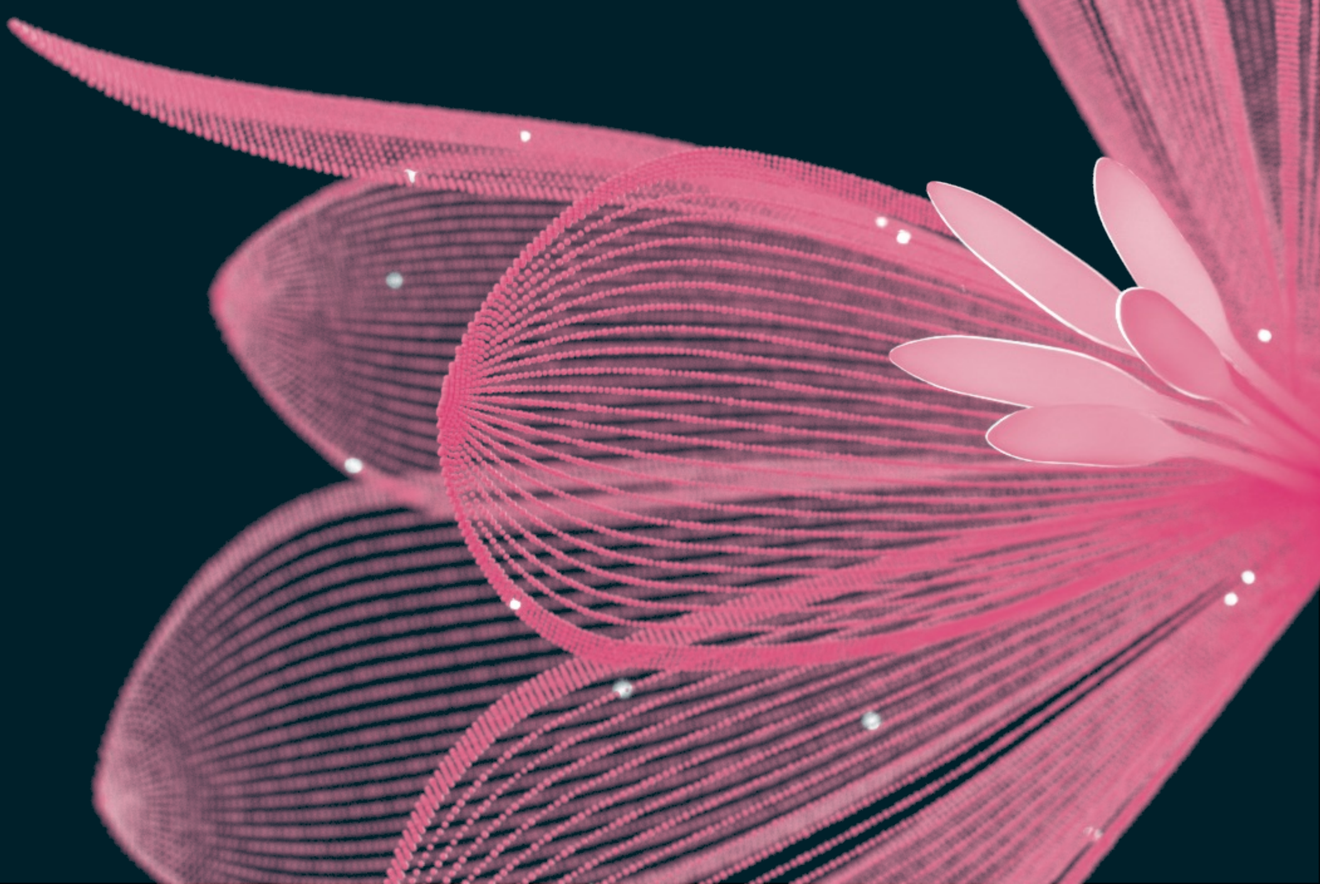

Microservice for AI-based Predictive Analytics

Make better and faster decisions by harnessing the power of AI in AVEVA PI with Amber. Amber detects equipment noncompliance earlier than any other predictive analytics tool, giving you more time to schedule maintenance, order parts, or change production settings. Production and maintenance teams can use Amber's native AVEVA PI plugin to better utilize existing data to lower operating costs and improve operating efficiency.

Amber's visionary approach to predictive analytics trains for the normal instead of the abnormal. By building its models on compliant operation instead of historical failure modes, Amber increases accuracy and accelerates time to deployment. With Amber, users can add dozens of tags to train an AI-based model in minutes without any manual model tuning or training, all within their AVEVA PI System.

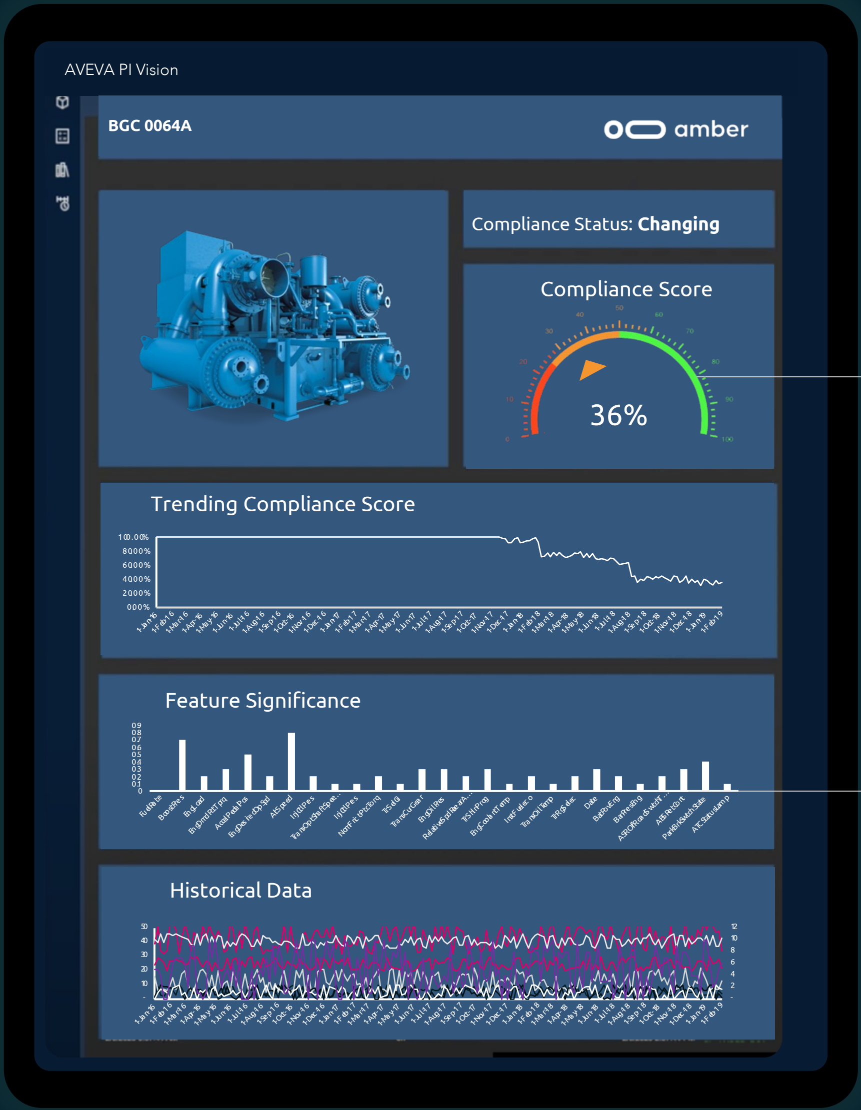

Compliance Score

Quantifies the degree of compliance exhibited by an asset.

Normal	51% – 100%
Changing	26% – 50%
Critical	0% – 25%

Feature Significance

Highlights the significance of each feature in contributing to a given Compliance Score.

CUSTOMER CASE STUDY

Optimized PM Schedule:

Linde Gas uses Amber to prevent $800k compressor failure

The world's largest industrial gas company supplies gas for a semiconductor manufacturer in Taiwan. Even with an on-site Level 3 vibration analyst, predicting failure has been difficult. Amber was integrated as a plugin to their AVEVA PI System to provide real-time insight into the health of each Atlas Copco compressor.

On January 29th, Amber noticed that the health of the asset was starting to change. On March 9th, the Compliance Score dropped below 25%, indicating that the compressor was now in a critical state.

"Amber alerted our staff members that they needed to stop and repair the compressor, which they did three months earlier than originally planned. That allowed them to see several cracks on their cooling bundle that, if left unrepaired, could have caused $800,000 in repair costs. Instead, they spent $60,000 to repair the cooling bundle and were up and running in 72 hours."

Paul Chen, System Integrator for AIONT, an Atlas Copco Partner

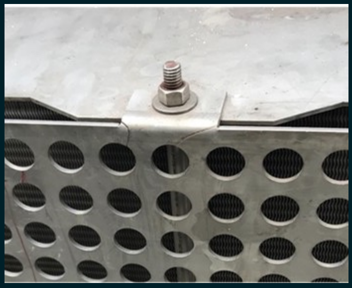

Date	Compliance Score	Event
Jan 6	98%	Training complete
Jan 29	50%	Changing asset health
Feb 11	40%	Continued degradation
Mar 9	24%	Critical asset health
Mar 13	19%	Maintenance conducted
Mar 16	95%	Repair complete

Read Full Case Study

Contents

01

Asking Reliability Engineers to Do the Impossible

In many cases, sites rely on maintenance and reliability engineers to establish appropriate thresholds that define the operational range (minimum and maximum) for sensors connected to specific assets. When a threshold is breached, the reliability team receives an alert and proceeds to examine the data to determine whether the alert represents a genuine issue or a false alarm. Critical assets often have lower thresholds, providing heightened sensitivity to potential breakdowns but leading to more frequent false alarms.

As the number of assets being monitored and the quantity of tags collected for each asset increases, reliability engineers may find themselves overwhelmed with configuring analytics for each asset and monitoring them thereafter. It is not uncommon for a single manufacturing site to have a substantial number of asset tags (sometimes as many as 100,000) in their PI System.

Even if there were enough reliability engineers available to analyze all the real-time data, asset noncompliance that is subtle and does not breach any established min-max thresholds would still be a challenge to detect.

1.1 The Challenge with Threshold-based Anomaly Detection

Manufacturing environments often consist of multiple pieces of equipment that are interrelated and cycle through various operating states. These operating states can be influenced by factors such as the product being produced, specific production settings, or seasonal variations.

Setting fixed thresholds for anomaly detection is a one-size-fits-all approach limited in its ability to adapt to the dynamic nature of assets in manufacturing environments. Only significant deviations from established thresholds trigger an alert, meaning that subtle anomalies or gradual changes in behavior often go unnoticed. Weeks of noncompliant behavior may be overlooked until the situation reaches a critical point, leading to potential disruptions in manufacturing processes and costly repairs.

To capture nuanced changes in asset behavior, more adaptive and intelligent approaches to anomaly detection are required. Models built on compliant operation allow for earlier identification of even subtle variations.

1.2 Normal Equipment Variation

- Different product types
- Varying input materials
- Various equipment operators
- Different weather conditions and times of year
- Changing production settings

02
A New Way of Anomaly Detection Detecting What Others Can't

To accurately assess an asset's health, an expert reliability engineer will ideally analyze multiple parameters simultaneously, taking into consideration past experiences and the relationships among various factors. When examining vibration patterns on a pump, for example, the engineer understands that the operating conditions, such as flow rate (50 GPM vs. 30 GPM) or the pumped material (water or slurry), can influence the vibration characteristics.

This comprehensive approach involves assessing what the pump is designed to do, how it is currently performing, and its measurement signatures. However, this level of sophistication is difficult to achieve using simple threshold-based or statistical anomaly detection.

Amber's AI-based predictive maintenance focuses on uncovering subtle relational anomalies that depart from historically compliant behavior, enabling identification of deviations in the interconnected relationships among parameters. Evaluation of these nuanced patterns allows Amber to detect potential issues early, even in complex and evolving operational conditions. This offers reliability teams a more comprehensive understanding of the asset's health and supports proactive maintenance and decision-making.

2.1 What is a Relational Anomaly?

What is an anomaly?
An anomaly is a measurement that deviates from what is standard and normal.

What is a threshold anomaly?
A threshold anomaly occurs when a single tag produces a value that is higher or lower than an expected or preset threshold.

What is a relational anomaly?
A relational anomaly occurs when two or more tags produce values that are normal when viewed individually but abnormal when viewed in comparison to each other.

Simple example
Suppose the following relationships were recorded for a pump:

Output (GPM)	Vibration
10	10
20	20
30	25

A reliability engineer sets a fixed Vibration threshold of 26. Now suppose the engineer also uses Amber to train a model of compliant operation with Output and Vibration as two features. The results might be as follows:

	Output (GPM)	Vibration	Threshold	Amber
Week 1	10	10	Normal	Normal
Week 2	18	20	Normal	Changing
Week 3	20	24	Normal	Critical
Week 4	8	14	Normal	Critical
Week 5	18	27	Critical	Critical

By looking at the relationship between Output and Vibration, Amber is able to detect an abnormal, never-before-seen relationship in Week 2. It would take three more weeks for the problem to become severe enough to be identified using standard threshold-based anomaly detection.

This example is very simple. In deployment, Amber can build models based on hundreds of asset tags that capture up to a thousand compliance relationships.

2.2 Detect what others can’t.

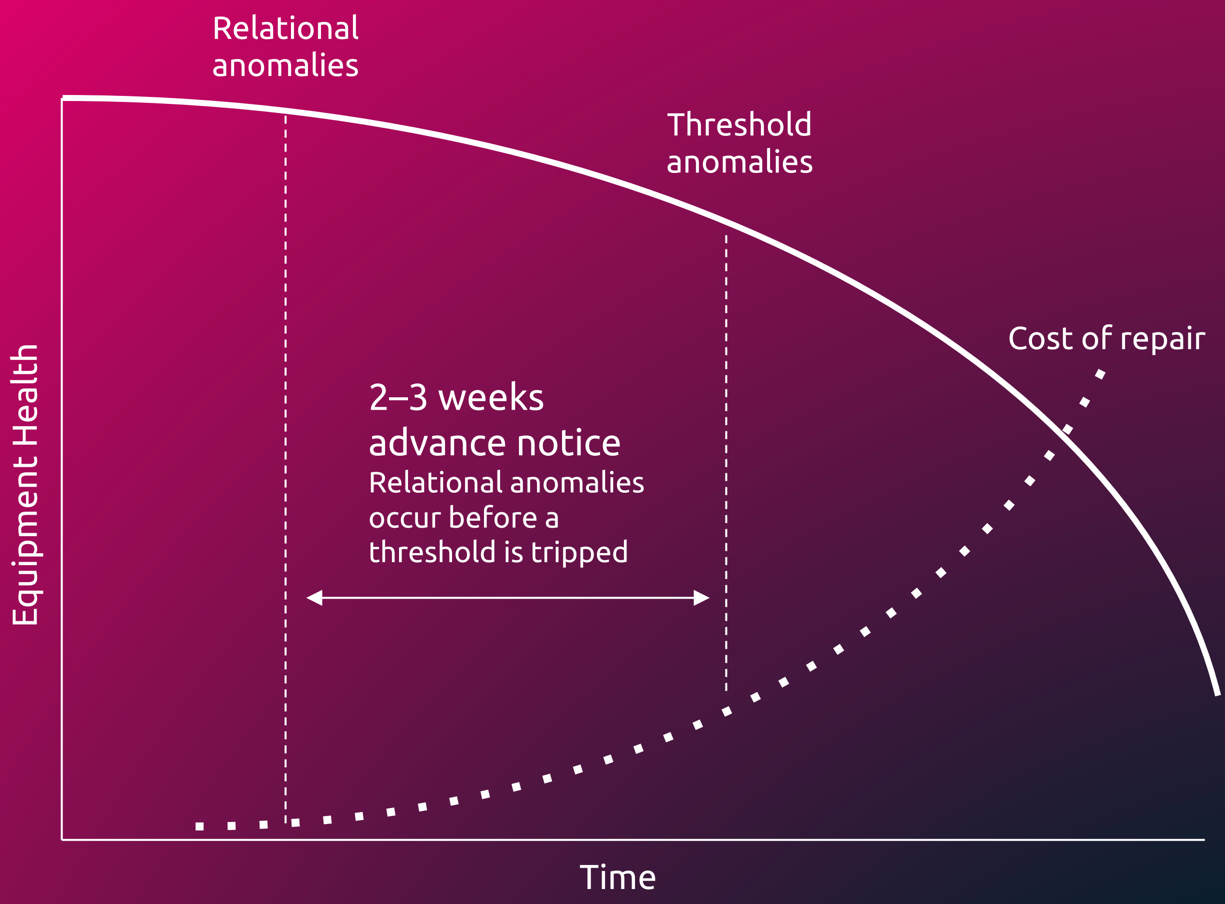

03
Path Toward Data-driven Decision Making

Predictive maintenance serves the crucial function of alerting team members when an asset begins to exhibit noncompliant behavior. Amber provides valuable detail about the extent of the noncompliance and the contributing factors behind the anomaly. This transformative insight enables team members to take informed actions to address potential failures.

Amber's **Compliance Score** quantifies the degree of noncompliance exhibited by the asset. It provides a numerical representation of the asset's deviation from expected behavior, allowing team members to gauge the severity of the anomaly. A lower Compliance Score suggests a more critical deviation, signaling the need for urgent attention and intervention.

Normal	51% – 100%
Changing	26% – 50%
Critical	0% – 25%

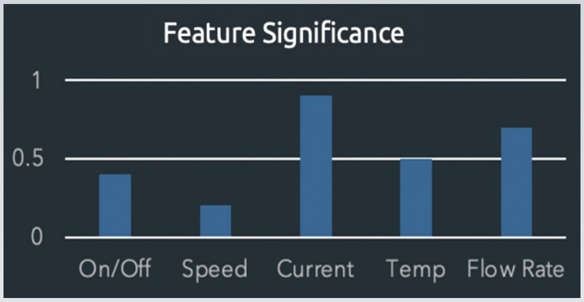

Feature Significance provides a rating of the likelihood that an individual feature is contributing to an anomaly. By identifying the specific parameters that play a crucial role in causing the abnormal behavior, Amber provides valuable information about the probable causes of the deviation.

3.1 Business Outcomes from Data-driven Decision Making

The synergy between Compliance Score and Feature Significance empowers reliability teams to assess the severity of a potential failure and make data-driven decisions regarding maintenance priorities, resource allocation, and mitigation strategies.

Operational Enhancement of Reliability Teams	Business Outcome
Perform maintenance before scheduled preventative maintenance (PM)	Avoid unplanned downtime or catastrophic damage
Avoid the need to schedule PM-based downtime of compliant assets	Increase the effectiveness and lower the cost of reliability programs
Order spare parts on an as-needed basis and with sufficient lead time	Reduce spare part inventory and supply
Shift staff resources from equipment monitoring to workorder execution and reliability program improvement	Reduce FTEs required and improve morale of program
Identify likely location and cause of asset issue before on-site inspection	Faster turnarounds and lower maintenance costs

3.2 Getting Ready for Amber's Predictive Maintenance

For a reliability program to reap the full benefit of PdM, it must be able to act effectively on the warnings received. For best results, the following should be present in a predictive maintenance program:

- A culture that seeks continuous improvement
- Implementation of PM schedules for all critical equipment
- Maintenance work orders that are closed in a timely fashion

3.3 Workflow with Amber

Issue Arises: A pump's inlet begins to experience clogging, leading to alterations in inlet pressure in relation to the other operational parameters of the pump.

Identify Issue and Root Cause: As the clogging intensifies, Amber's Compliance Score steadily decreases, dropping to 40%. Feature Significance highlights unusual correlations among flow rate, inlet pressure, power, and temperature.

Diagnosis and Work Order Generation: Prompted by notifications, staff suspect debris in the suction lines. A work order is generated for investigation and resolution.

Resolve Issue: The staff members clean the suction lines and reestablish the connection. As a result, Amber's Compliance Score rebounds to a level indicating that the issue has been successfully resolved.

04

How to Get Started with Amber in AVEVA PI

Amber's pre-built plugins simplify the process for AVEVA PI users to create numerous models within a few days. The process involves four steps:

1. Connect
2. Configure
3. Train
4. Monitor

Amber employs an unsupervised anomaly detection approach, eliminating the need for manual model training, fine-tuning, or data labeling. All Amber needs to train a model is data from an asset that accurately represents its typical operational behavior. Once operational, Amber's results can be seamlessly visualized in PI Vision or any other user interface.

4.1 Estimated Hours for Project Completion

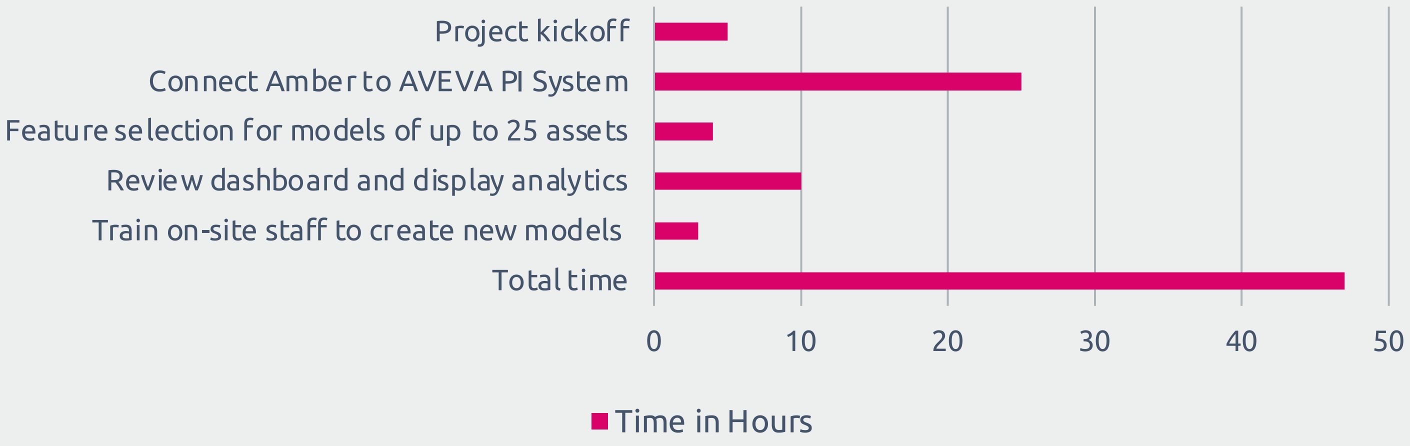

Step 1. Connect Amber to AVEVA PI

Amber seamlessly integrates with any existing data management platform, offering deployment options both on-premise and in the cloud. Subsequently, the analytical insights provided by Amber are visualized within the pre-existing user interface, which could include platforms like PI Vision, Power BI, and more.

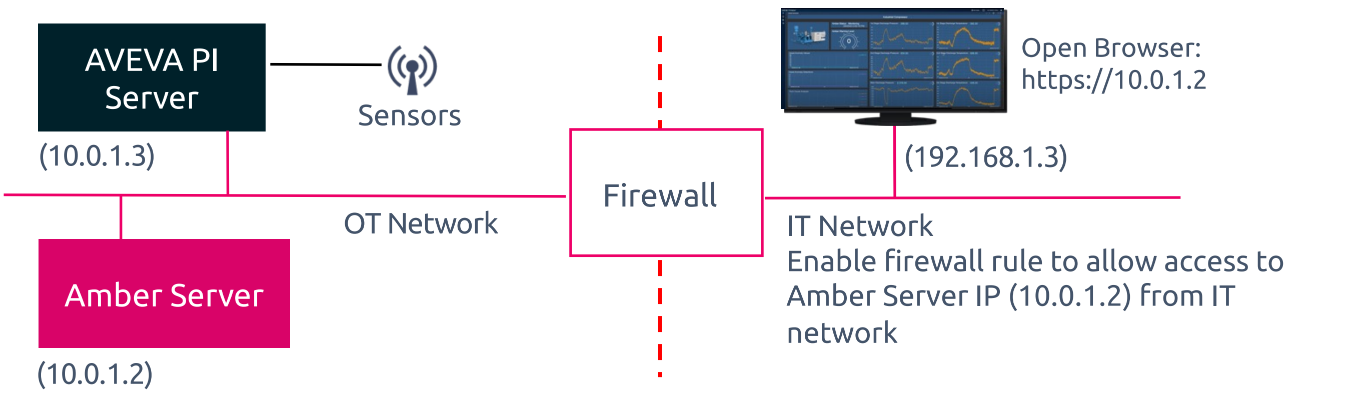

Technical Requirements

- Existing AVEVA PI Server
- PI Web API enabled
- Ethernet connection on the same network as PI server (OT network)
- Username and password credentials for PI Server
- Known IP address and domain name of AVEVA PI Server

Download Integration Document

Step 2. Configure Amber

After installation, users gain access to the Amber configuration page, where they can initiate the process of creating a new model. This involves selecting tags from the list within the PI Asset Framework to incorporate into the model's setup. It is important to note that each model has the capacity to encompass a maximum of 500 tags.

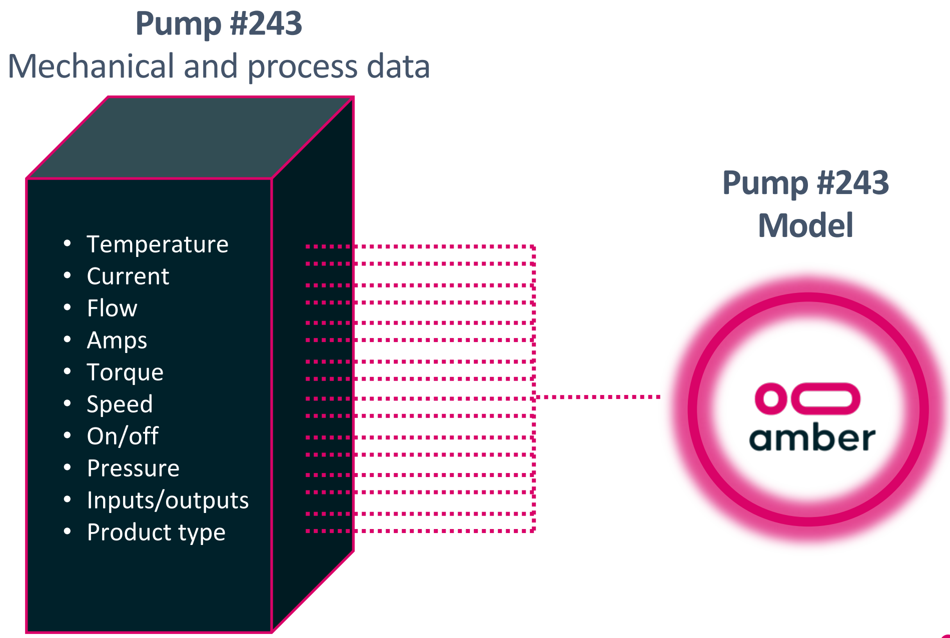

Types of Tags to Include

- Production inputs and outputs
- Equipment settings
- On/off
- Temperature
- Current
- Flow
- Speed
- Pressure
- Displacement measurements
- Measurements of manufactured widgets, such as size, weight, or other dimensions

After choosing the appropriate tags, proceed to pick a historical data range spanning 1-6 months for training Amber. While selecting this dataset, it is crucial to confirm that:

1. The asset is functioning in a compliant state.
2. The dataset is sufficiently extensive to encompass all typical operational states the asset regularly experiences.

In the realm of heavy industrial equipment, flawless operation is an unattainable ideal. The pivotal criterion for selecting an asset in a compliant state hinges on confirming the absence of any indications of smoke, unusual noises, or deviations from its customary behavior. In simpler terms, if you are satisfied with how the asset is performing, it meets the criteria for Amber's training.

Certain assets, like the motors in a collaborative robot (cobot), exhibit a highly repetitive operational pattern recurs every few seconds. Conversely, equipment such as a pump in a specialty chemical manufacturing setup may feature numerous distinct operating modes, contingent upon the manufacturing stage and the particular prothatduct in production. In both scenarios, the quantity of data required to represent typical operation varies significantly. Therefore, the training dataset should be sufficiently extensive to encompass multiple iterations of all standard operating conditions for the respective asset. As a rule of thumb, the more data the better.

For example, when training a pump that operates with freshwater, saltwater, and a silty slurry, the training dataset should encompass all three scenarios. Failing to include such diverse scenarios might lead to false alerts during Amber's monitoring phase. An untrained scenario, such as not training for silty slurry in this case, could result in a warning when the condition is encountered for the first-time during monitoring. In such an event, learning can be reenabled to include the new operational state.

Considerations for Data Selection

- Include all normal equipment variation including production settings and environment factors
- Equipment must be in a “fine” operating condition (no smoke, fire, or major problems)

Step 3. Train Amber with Auto-learning

After a sequence of features and a historical data timeframe has been selected for an asset, Amber is prepared to initiate auto-learning. With the press of a button, Amber undertakes the task of segmenting the training data into distinct groups based on their similarities and differences. At this point, no more user intervention is required. The system dynamically adjusts the quantity of clusters formed to encompass all typical variations inherent to the specific asset.

Sensor Fusion Vectors

Amber compiles individual data rows into a consolidated vector, termed a "sensor fusion vector." This vector is essentially a sequence of values that captures an asset's state at a particular point in time.

Freq	RMS	Pos	Temp	Vib	Press
52	119	134	74	10	33
58	116	139	74	12	35
62	100	152	75	11	39

Sensor Fusion Vector #23

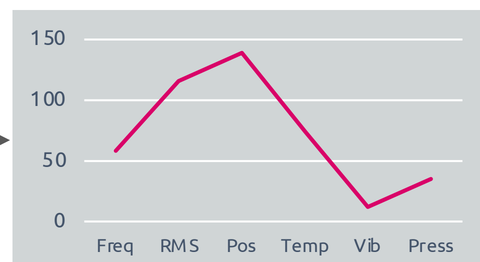

Cluster Creation

Each sensor fusion vector is sorted into a cluster with other similar vectors. Should Amber encounter a sensor fusion vector that significantly differs from existing clusters, it will establish a new cluster. Empowered by Nano, Amber has the ability to form an endless number of clusters. Consequently, the quantity of clusters within the model is dictated by the training data unique to each asset.

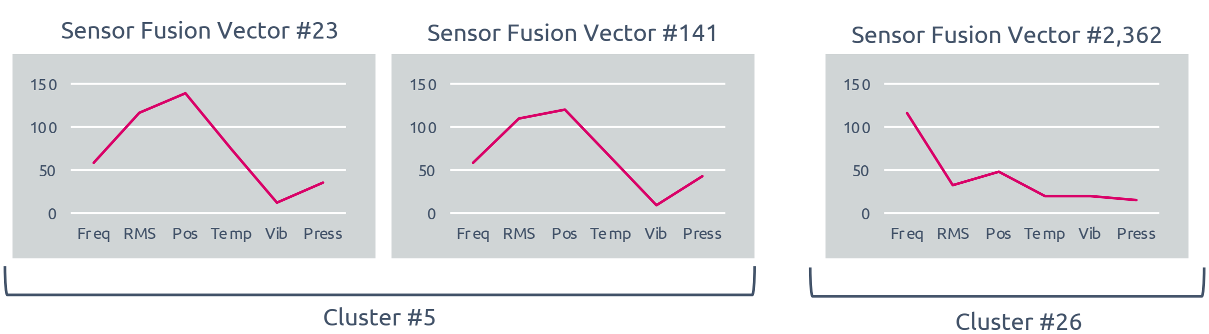

Sensor fusion vectors #23 and #141 are similar and are placed in the same cluster, whereas sensor fusion vector #2,362 is very different and so is placed in a separate cluster

Learning Curve

As the number of samples processed by Amber increases, the creation of new clusters diminishes, indicating that the system comprehends the intricacies of the asset's operational patterns. This trend is reflected in the learning curve, depicting the correlation between processed sensor fusion vectors and the generation of fresh clusters. A flattening of the learning curve signifies the maturation of the model.

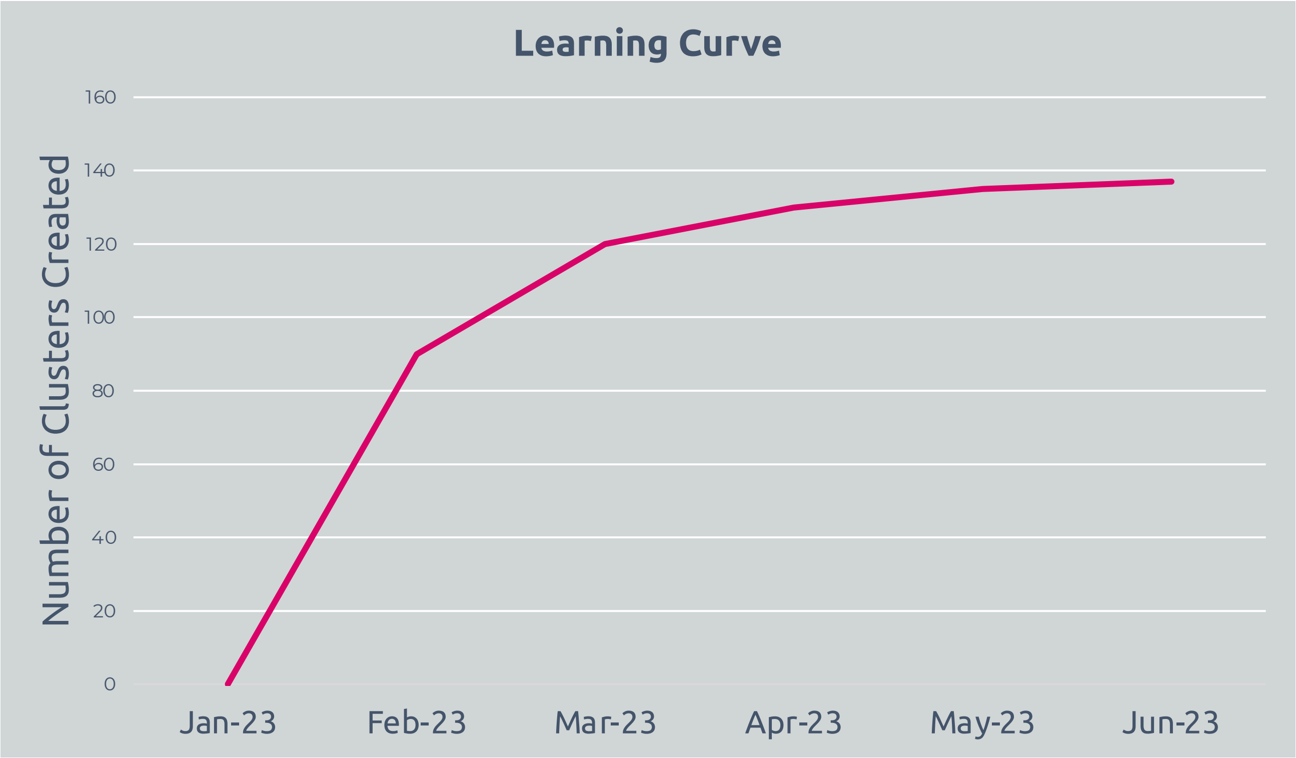

In January and February, Amber is exposed to new sensor relationships that it hasn't seen before. During this time period, 87 new clusters are created. From March to June, most operating states of the asset are familiar to Amber and only 19 new clusters are created.

Step 4. Commence Monitoring

After training concludes, Amber seamlessly shifts into monitoring mode. Throughout this phase, streaming data is assessed against the model established during training. For every sample, a Compliance Score and Feature Significance rating are generated. Samples closely resembling the model yield higher Compliance Scores, whereas significantly deviating samples result in lower Compliance Scores.

Compliance Score

In the example below, the Compliance Score (CS) is 100% for several weeks. In December, CS begins to show some variability, but it remains above 50%, indicating that the asset's behavior continues to align with its behavior during training. CS drops below 50% for the first time in September, suggesting that the asset is transitioning into an uncharted operational state. A CS of 25% or lower signifies that the asset is currently in a new state.

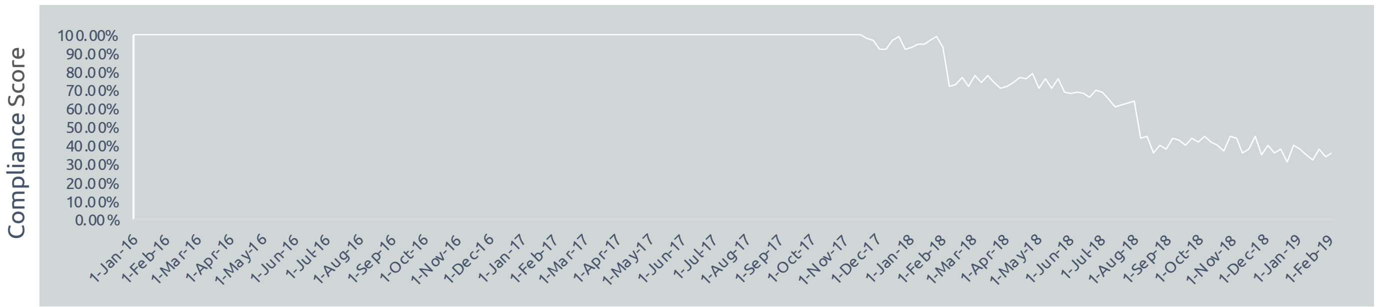

Feature Significance

Feature Significance (FS), rated on a scale from 0–1, indicates the likely effect of each tagged feature on a particular sample's CS. The three to four features with the highest FS values are those that the reliability team should investigate as being responsible for the change in asset compliance.

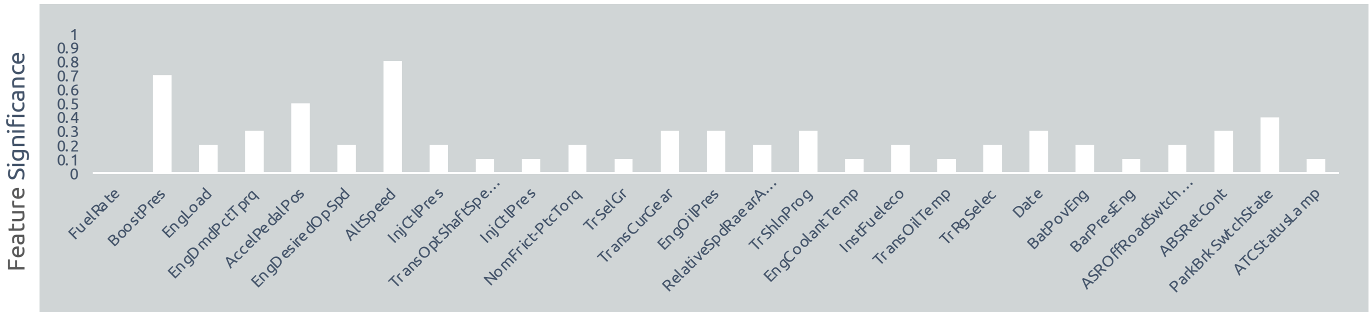

The New Normal in Anomaly Detection

Explore the demo and discover anomaly detection like you've never seen it before.

Amber integrated into PI Vision

Watch a Demo