



The 5 Step Roadmap To IoT-Based Predictive Maintenance

How to leverage sensor data, predictive
analytics & machine learning for more
intelligent maintenance

Published by XMPRO, Inc.

Copyright © 2016 by XMPRO, Inc. All rights reserved. Digitally printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior permission of the publisher.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is provided with the understanding that neither the author nor the publisher is engaged in rendering legal, accounting, or other professional service. If legal advice or other expert assistance is required, the services of a competent professional person should be sought.

Digital Copyright XMPRO, Inc. © 2016

External Data: Some of the data included in this publication has been extracted from external sources, therefore all such data included herein is the outcome of the analysis by XMPRO staff of data from aforesaid sources. XMPRO claims no right on this data and has done its utmost best to reference data within this document to its original respective Copyright owners. However, the conclusions provided within this document are the point of view of XMPRO.



48446752729023365

TABLE OF CONTENTS

| | |
|--|-----------|
| The Evolution of Maintenance | 4 |
| A Brief History | 4 |
| Where we are now | 6 |
| 1. Identify Assets | 8 |
| Asset Evaluation Criteria for Predictive Maintenance | 8 |
| Analyze assets using RCM | 8 |
| 2. Gather Data | 10 |
| Install Sensors | 10 |
| Prepare Data | 11 |
| Feature Selection | 12 |
| 3. Design Predictive Model | 13 |
| Choose an Algorithm | 13 |
| Train the Model | 14 |
| Evaluate The Model | 14 |
| Methods to Measure Model Performance | 15 |
| Which method works for PdM? | 15 |
| Deploy the model into Production | 16 |
| Design the Eventstream in XMPRO IOM | 16 |
| 4. Manage Work | 18 |
| Why it matters | 18 |
| Assigning work | 19 |
| Resource Scheduling | 19 |
| Discussions | 19 |
| Notifications & Escalations | 20 |
| Field Inspections | 20 |
| 5. Learn & Act Smarter | 20 |
| Save Adaptive Case Action Sequences As Templates | 20 |
| Use Predictive Models to Predict Best Next Actions | 21 |
| Scale PdM Program Across Multiple Asset Types | 22 |
| Move from Predictive to Prescriptive Analytics | 22 |

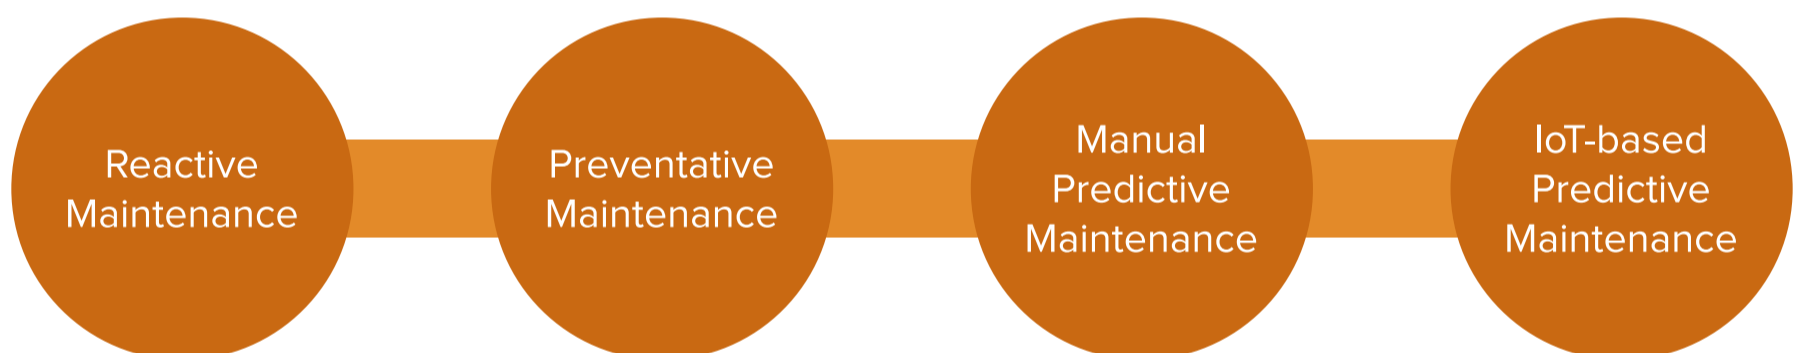
THE EVOLUTION OF MAINTENANCE

*“Change is inevitable -
except from a vending machine.”*

- Robert C. Gallagher

Since the invention of the first machine, man has had to deal with maintenance and repairs. Thankfully, as our technology has become more sophisticated, so has our approach to maintaining it.

A BRIEF HISTORY



It started out with reactive maintenance, operating on a ‘fix it when it breaks’ model. This approach has very obvious flaws, especially when used in large-scale production environments. We soon realized that catastrophic failures weren’t good for business.

This gave way to a new approach. We decided not to wait until something breaks, but rather to perform preventative maintenance at fixed time intervals to stop these failures from occurring. While preventative maintenance might have offered some benefits, it didn’t solve truly the problems we were facing.

Here are a few reasons why:

- It creates unnecessary work, by sending technicians to fix things that aren't broken.
- It's labour intensive, requiring a large number of technicians to manage a company's assets.
- Preventative maintenance often doesn't uncover critical issues that arise between maintenance intervals.
- According to NASA, failure patterns that are age related only apply to 18% of assets.¹

The combined cost of excess maintenance and lost productivity in the US has been estimated at \$740B.²

Then things changed again. Companies started to take advantage of the benefits predictive maintenance (PdM) had to offer. The ultimate goal of implementing predictive maintenance was to prevent unplanned reactive maintenance due to failure and eliminate the costs of doing unnecessary preventative maintenance. This approach keeps the maintenance frequency as low as possible while still preventing equipment failure.

According to a survey³ by Plant Services, 7 out of 10 respondents have a predictive maintenance program in place, with the most common technologies being vibration analysis, infrared imaging, and oil analysis. We've labelled PdM programs using these techniques 'Manual Predictive Maintenance'.

1 Proactive Asset Management

2 Simplifying Predictive Maintenance

3 Is your PdM Program as good as it could be?

WHERE WE ARE NOW

The next step in the evolution of maintenance is to move to a more scalable IoT-based predictive maintenance program. This approach to maintenance used to be out of reach for most organizations, because of the high cost of computational memory and the lack of scalable platforms.

The Internet of Things has opened up unprecedented new opportunities for companies in industrial verticals like oil and gas, utilities, manufacturing, and aerospace. Data created by industrial equipment has potential that goes far beyond the hype of wearable consumer electronics. It is no surprise that industry leaders are looking to these new technologies to drive innovation and efficiency.

In a recent survey by GE to executives in industrial verticals, 76% of respondents said that they plan on increasing their spending in Big Data analytics. This increase in spending is driven largely by the fear of their competitors gaining market share at their expense. Another driver is the belief that Big Data Analytics has the power to change the competitive landscape they operate in.⁴

Using predictive analytics to drive maintenance is what we've labelled "IoT-based predictive maintenance". It allows you to scale your predictive program by leveraging IoT data and predictive models, without incurring massive costs or scrambling to find more qualified employees.

Data is the new competitive asset and those who can glean the most insight from their data will have the upper hand when making decisions. Just like any other asset, its value is tied to the expectation of receiving benefits in the future.

Many companies are data rich and information poor at this stage. They have sensors on their equipment and they're generating masses of timestamped temperature and pressure readings, but they need to take things a step further to truly start seeing the potential Big Data has to offer.

4 Industrial Internet Insights Report for 2015

Companies that get the highest yield of insights from their data will rise above the rest in the years to come. As your analytics program becomes more mature you'll be able to glean more insights and progressively improve your decision-making capabilities, allowing you to get more value out of the data you already have.

We've found that the lowest hanging fruit with regard to Big Data for industrial companies lies in using predictive analytics to enable an IoT-based predictive maintenance program. In this ebook, we're going to show you how to get started with the 5 Step Roadmap to IoT-Based Predictive Maintenance.

5 STEP ROADMAP TO IOT-BASED PREDICTIVE MAINTENANCE



Identify Assets



Gather Data



Design Predictive Model



Manage Work



Learn & Act Smarter





1. IDENTIFY ASSETS

Now that you've decided to implement an IoT-based predictive maintenance program, you need to determine which assets are a good fit. Here is a simple framework to evaluate which assets would be best suited to a PdM program.

ASSET EVALUATION CRITERIA FOR PREDICTIVE MAINTENANCE

- Does the asset have a critical operating function?
- Is there significant cost associated with repairing the asset after a failure?
- Are there safety risks associated with asset failure?
- Is there significant cost associated with performing preventative maintenance on the asset?
- Can you predict failures in a cost effective way through monitoring and a predictive model?

ANALYZE ASSETS USING RCM

Next, we are going to use reliability centered maintenance to perform an analysis on each asset we've chosen to include in the PdM program. RCM provides a proactive assessment of the likely failure modes for an asset and helps to develop specific tasks to prevent reoccurring failures.

RCM was developed in the 1960s for the airline industry and has since been used in the space, defense and nuclear industries, where functional failures have the potential to cause extreme environmental impact and loss of life. While performing a rigorous RCM analysis is appropriate for those industries, we're going to use a more streamlined version to analyze our equipment.

Combining the RCM assessment with predictive analytics not only gives you a clearer indication of the status of your equipment but also enables you to efficiently manage the work that results from those insights.

Implementing a full-scale RCM program and predetermining the tasks for each possible failure mode for each piece of equipment can be a very daunting, expensive and time-consuming process. If that approach doesn't appeal to you or it's just not within your budget, we'll show you how to get around that using our Adaptive Case Management in Step 4 of the roadmap.

Pairing an RCM program with predictive analytics will allow you improve the answers to your 7 questions over time, enabling even better data-based decision making.

There are 7 questions that need to be asked when using the Reliability Centered Maintenance strategy:

1. What are the functions and desired performance standards of each asset?
2. How can each asset fail to fulfill its functions?
3. What are the failure modes for each functional failure?
4. What causes each of the failure modes?
5. What are the consequences of each failure?
6. What can and/or should be done to predict or prevent each failure?
7. What should be done if a suitable proactive task cannot be determined?

A failure mode is any event which is reasonably likely to cause a functional failure in the asset. This can be due to deterioration, human errors, design faults and poor procedures.⁵

Failure modes can be divided into 3 groups:

1. Failures that have occurred before on same or similar asset
2. Failures that have been successfully prevented by existing maintenance
3. Any failures outside those 2 categories which are considered a real possibility in the future



2. GATHER DATA

INSTALL SENSORS

The first step in setting up your predictive maintenance infrastructure is to install sensors on your assets. Subsequently, you need to either set up a gateway to the Internet to stream data from those sensors or you can send the sensor data to your Plant Information system.

Following the RCM assessment, you can decide which sensors you want to install on your equipment based on the causes of failure modes you want to predict and the indicators of those failure modes. To decide on the infrastructure required to form the foundation of your predictive maintenance program, you need to ask the following 2 questions:

1. What do you want to measure?

Examples are:

- Pressure
- Vibration
- Temperature
- Voltage
- Flow

2. What do you want to do with your data?

- Store in an external database
- Send to a cloud service
- Store in a Plant Information system

PREPARE DATA

Why is it important?

Quality data is the lifeblood of a quality model. Even though sensors provide a higher level of data accuracy than manual entry would, it still doesn't make for a completely clean dataset right off the bat.

Like Eric King says, "Many people assume that modeling greatness is wrapped up in the ability to build a better algorithm. Building a better algorithm is like building a faster rocket ship. It's great, as long as you're pointing the ship in the right direction. Nobody wants to move faster if they're going in the wrong direction, but this is precisely what happens in many organizations. Data preparation is a big part of keeping that rocket pointed in the right direction."⁶

According to The Data Warehousing Institute (TDWI), the unusable nature of dirty data (a database record that contains errors) costs organizations up to \$600 billion each year.⁷

The quality of your data is going to affect your model's performance. Model performance is going to affect the maintenance schedule you follow. Your maintenance schedule is going to affect both maintenance and production costs. So while data preparation might not sound glamorous, it is of key importance when implementing a PdM program.

What tools should I use?

Data preparation can be done either in XMPro or the Azure Machine Learning Studio.

- XMPro's Stream Designer includes a set of extensible/pluggable transformations to transform your data. Data can then be sent to the Azure model after preparation.
- If your data is already in a tabular format you can also perform data pre-processing in the Azure ML Studio.

6 Data Preparation: A Vital Step in Predictive Analytics

7 Data Quality & The Bottom Line

Typical Data Quality Issues:

- Incomplete: Data lacks attributes or contains missing values.
- Noisy: Data contains erroneous records or outliers.
- Inconsistent: Data contains conflicting records or discrepancies.

Common data preparation tasks:

- Data cleaning: Fill in missing values, detect and remove noisy data and outliers.
- Data transformation: Normalize data to reduce dimensions and noise.
- Data reduction: Sample data records or attributes for easier data handling.
- Data discretization: Convert continuous attributes to categorical attributes for ease of use with certain machine learning methods.
- Text cleaning: remove embedded characters which may cause data misalignment.

For more detailed information go to [Learning Guide: Advanced data processing in Azure](#).

FEATURE SELECTION

Feature selection is the process of identifying the most influential variables in your dataset. It is crucial to understand the role played by each feature, especially when working with large datasets. Some variables will be better for predicting equipment failure than others. Reducing the features in your dataset will prevent the algorithm from having a high variance. Less features also reduce the computational power needed to run the model. This allows for faster experiments.

Feature selection is often a manual process that requires expert domain knowledge. Experienced members of your engineering and maintenance team will be of great value if you choose to go this route.

Another way to do feature selection is to use Azure's Filter Based Feature Selection module. This module allows you to apply different statistical tests to your dataset, based on the type of data and task. The feature columns with the best scores are then used to build the predictive model.



3. DESIGN PREDICTIVE MODEL

CHOOSE AN ALGORITHM

The business problem you are trying to solve is going to dictate which algorithm is most appropriate.⁸ Now you need to ask “what do we want to predict?”

Predicting if it will fail - Use Binary Classification

Binary Classification attempts to predict whether an asset falls into one of two classes. The classes are mutually exclusive and can be called ‘will fail’ and ‘will not fail’ for example.

Azure Machine learning gives you multiple two-class classification algorithms to choose from including:

- Two-class logistical regression
- Two-class boosted decision trees
- Two-class decision forest
- Two-class neural networks

Predicting when it will fail - Use Survival Analysis

Survival Analysis is used to predict the time until the occurrence of an event of interest. In the case of predictive maintenance we want to predict the time until an asset failure.⁹

8 Predictive Analytics with Microsoft Azure Machine Learning

9 What is Survival Analysis?

For more information on choosing the right algorithm for your predictive model have a look at this cheat sheet - [Choosing an Algorithm Cheat Sheet](#)

TRAIN THE MODEL

The next phase of the experiment is to split your data into two samples. Your training sample is going to comprise 70% of the initial dataset. The remaining 30% is going to be allocated to the test sample. You can now use the 'Train Model' module in the Azure Studio to run your experiment and train the model.

EVALUATE THE MODEL

Compare Algorithms

Azure gives you the ability to score and evaluate multiple algorithms and find the most accurate model for your dataset. Algorithms are compared using the same training and testing data samples.

Measure Model Performance

Azure has provided multiple methods to measure your model's performance.

Why do you need multiple methods? Because the ratio between positive values and negative values in the data, otherwise known as the class distribution, isn't always 50/50.


The cost of wrong predictions isn't equal either. False positives and false negatives have different costs associated with them. In the case of a false positive in predictive maintenance, the model predicts a failure when there isn't going to be one. This could lead to incurring the cost to get a maintenance person to inspect the piece of equipment and find that there is nothing wrong. Or it could mean replacing a part that didn't need to be replaced.

Compare that to the cost associated with a false negative, where an asset failed and had to be fixed reactively. Assuming the asset meets the criteria we used in Step 1, false negatives could have very high cost and safety risks associated with them. It makes sense to optimize the model to predict **less false negatives** when critical operating equipment is involved.

METHODS TO MEASURE MODEL PERFORMANCE

Receiver Operating Characteristic Curve

The ROC curve gives you a visual representation of your binary classification model's performance. It has True Positive Rate on the y-axis and False Positive Rate on the x-axis.



Precision ¹⁰

Precision gives you the ratio of correct positive observations.
 $\text{True Positives} / (\text{True Positives} + \text{False Positives})$.

Recall

Recall gives you the ratio of correctly predicted positive events.
 $\text{True Positives} / (\text{True Positives} + \text{False Negatives})$.

F1 Score

The F1 score gives you the weighted average of both Precision and Recall.
 $2 * (\text{Recall} * \text{Precision}) / (\text{Recall} + \text{Precision})$

WHICH METHOD WORKS FOR PDM?

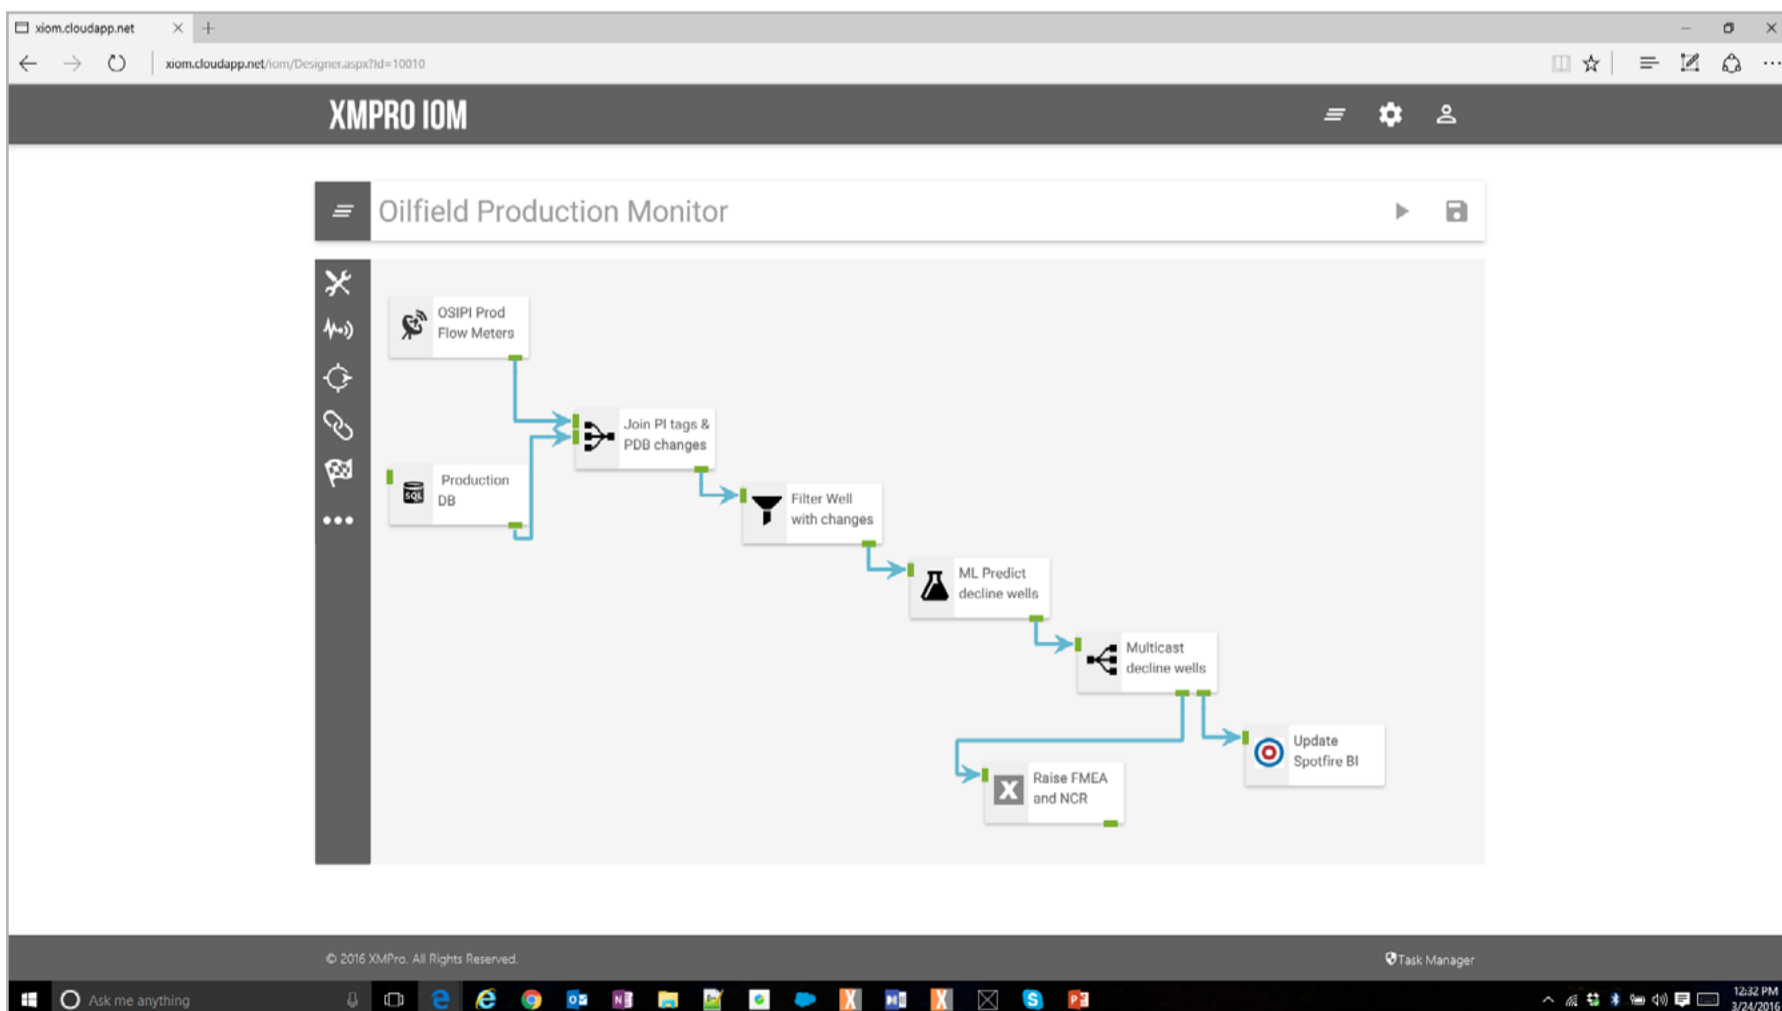
In the case of predictive maintenance where the cost of a false negative (failure happens and wasn't predicted) is higher than a false positive (fixing something that wasn't broken), it is best to focus on the Recall method to evaluate the model performance.

DEPLOY THE MODEL INTO PRODUCTION

Now that you've completed the model you can save it and remove the training modules in Azure. Moving your model into production is done by publishing it to a new web service. Here you'll receive the APIs to call your model as a web service and use it in XMPro's Stream Designer.

DESIGN THE STREAM IN XMPRO

XMPro's Stream Designer gives you the ability to configure the entire predictive maintenance stream for an asset using a drag and drop interface. This includes everything from setting up the connection to the sensor data or PI system, to running the predictive model, sending notifications and initiating new actions.



Configure Plant Information System Queue

We're going to start by setting up a connection to our Plant Information System, which contains all of the data from our sensors. XMPRO seamlessly connects to the external systems without requiring any coding. It also has a library of standard connectors, which includes a connector to the popular Plant Information System by OSIsoft.

Set up Buffer

Next, we are going to set up a buffer to listen for a specific type of asset, like a transformer or a centrifugal pump.

Set up Filters

We can then set up a threshold for a specific feature that would indicate a failure, like a temperature reading above 400° or a flow rate of 0. You can also set up a notification that will immediately send an SMS alert to notify relevant parties about a failure.

Add Predictive Model

Here you will set up the connection to your Azure predictive model by using the API to call the web service you published earlier. In this instance, we are going to use Azure's data preparation tools.

Set up Actions

XMPRO lets you set up a response based on the prediction by the model. If the model predicts that a failure is likely to happen it can initiate an XMPRO case, which sends the information to an engineer. The engineer can then decide which steps should be taken for that particular situation.



4. MANAGE WORK

WHY IT MATTERS

The value of predictive maintenance lies in doing the maintenance work to prevent a failure **before it happens**. Alerts and notifications aren't enough to ensure that gets done, especially if you are implementing a large-scale PdM program across multiple types of assets. Tasks are bound to slip through the cracks if they aren't managed, assigned and tracked. Even though predictive maintenance inherently creates a more predictable maintenance program, you still need to be able to make decisions on a case-by-case basis. Every possible failure situation can't be planned for every asset.

If the model predicts a failure with a predetermined response task, you can kick off an automated workflow in XMPro.

The problem with most predictive maintenance programs is that the tasks associated with performing the maintenance don't allow for knowledge workers to make decisions. The system just creates an automated work order and doesn't provide the flexibility for a knowledge worker, like an engineer, to determine appropriate actions on a case-by-case basis.

In many cases there are unplanned events that happen and the tasks to respond to that situation aren't clear cut. This is especially true if you haven't moved into a full Reliability Centered Maintenance program. There are many times where you need to rely on an engineer or expert to look at the data and decide what should be done next.

In cases like this, you can kick off a case in XMPro, where the engineer can look at the data and the prediction and make a decision on which steps to take to resolve that particular situation. They then become the case manager and can create and assign tasks to technicians. In order for a PdM program to succeed you need to be able to capture the knowledge and experience from your workers and leverage it for more intelligent operations. This also allows you to improve your answers to the 7 RCM questions based on the new data created by your predictive maintenance program.

ASSIGNING WORK

After the engineer has decided which steps are appropriate for a particular case they can create tasks and assign them to specific technicians with deadlines. This ensures that the work gets done to prevent the asset failure.

Another important aspect of managing the work associated with a PdM program is the need for a fully visible audit trail related to the case. XMPro gives you a full audit trail of actions, files, and discussions associated with a case and allows you to easily create reports.

RESOURCE SCHEDULING

In order to manage work created from predictions, you need to invest in a tool that allows you to schedule the technician, asset and maintenance facility or equipment. XMPro's platform comes with an easy to use drag & drop scheduler that allows you to assign resources (both people and assets) to tasks in a visual interface. It also gives you the option to auto-schedule tasks when presented with a conflict in the schedule.

DISCUSSIONS

Collaboration is an essential part of making decisions as a knowledge worker. You want to empower your engineers and maintenance managers to consult with their colleagues in order to make the best decision for the situation. XMPro's discussions and crowd questions facilitate collaboration by embedding social-style messaging in the context of cases and processes. This gives your engineers the ability to discuss a task, ask advice and understand context without using email. All of the discussions are linked to the case and can be viewed as part of the audit trail to help you understand why a decision was made.

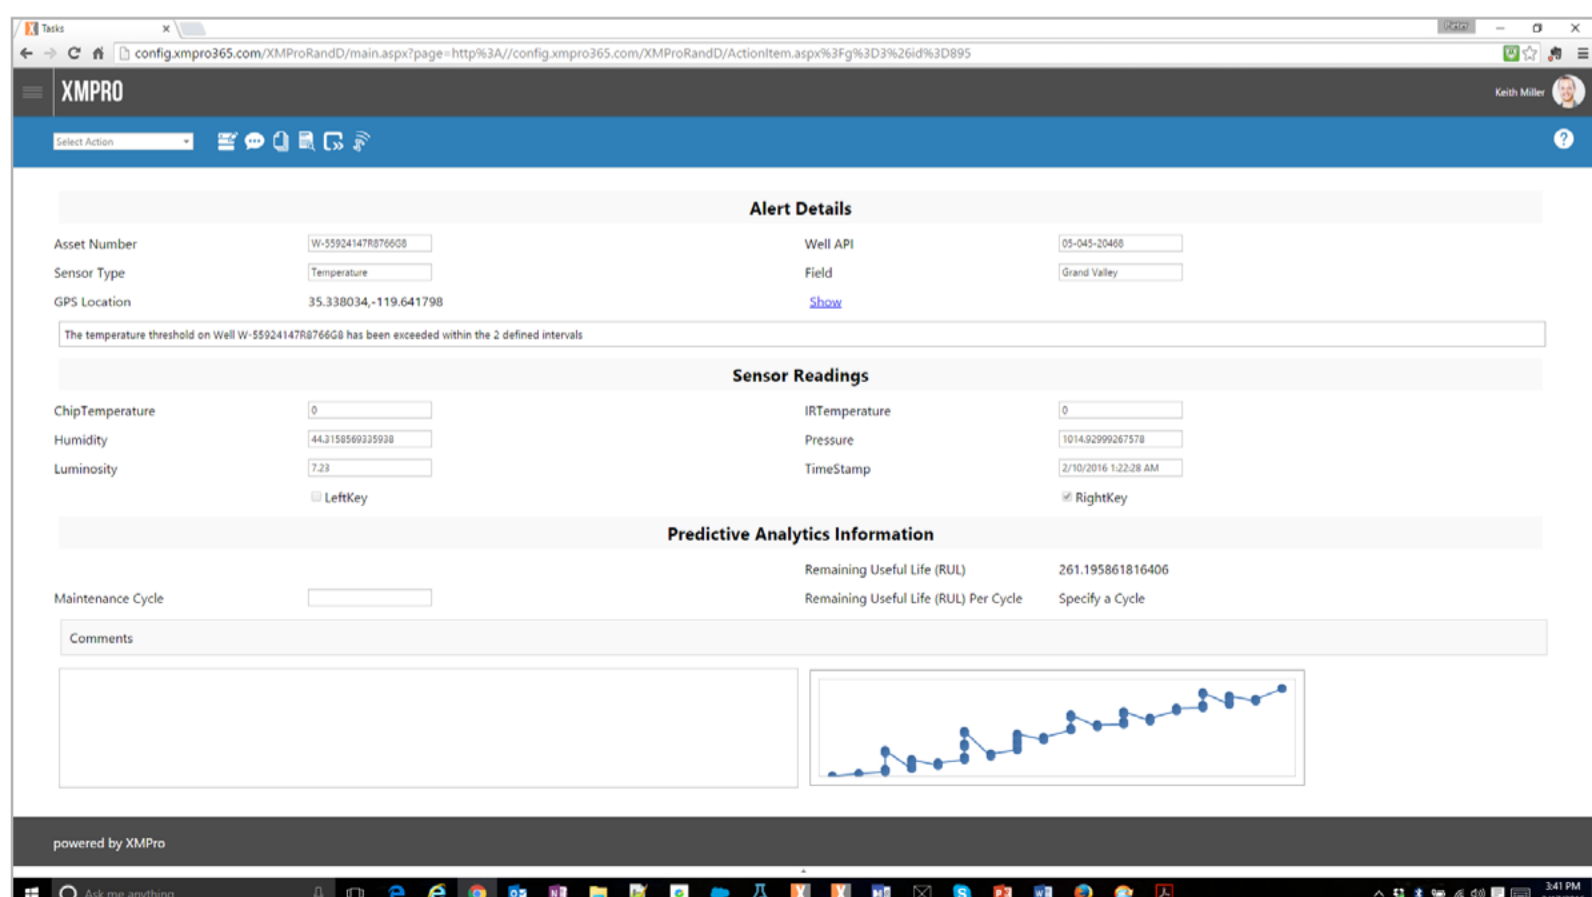
NOTIFICATIONS & ESCALATIONS

While notifications aren't enough to ensure a job gets done in time, they do assist with creating awareness about the situation. Your technicians can automatically receive notifications of new tasks via email and SMS.

It is also important to have an escalation procedure in place, especially when the task relates to the predicted failure of a critical asset. The work needs to get done in time. If a task isn't completed within a specified period of time the technician will receive another notification as a reminder. A task can automatically escalate to a designated person, like a maintenance manager, if it has not been attended to after multiple notifications.

FIELD INSPECTIONS

Your technicians need to be able to manage work from wherever they are. XMPro's mobile interfaces have fully featured task management capabilities. Technicians can also add photos, audio and video files to inspection forms from their mobile device. This means no more clipboards and paper checklists. Every action and file get added to the case, which can then be referenced in the future.





5. LEARN & ACT SMARTER

SAVE ADAPTIVE CASE ACTION SEQUENCES AS TEMPLATES

A key feature of XMPro's Case Manager is the ability to save a sequence of actions as a template, in order to capture the decision-making knowledge from the engineer. And create a repeatable process for when a similar situation presents itself in the future. This allows you to build up a knowledge base to aid decision making in the future, making your data a strategic competitive asset.

USE PREDICTIVE MODELS TO PREDICT BEST NEXT ACTIONS

Once you have a solid IoT-based PdM program in place, you can start to use machine learning and predictive models to predict the best actions to take in each situation.

XMPro's unique event-based architecture lends itself to process mining, which gives it the ability to analyze process information to look for patterns of work, predict the outcome of the process and suggest process steps that will guide actions towards desired outcomes. This gives you insight into the which actions created the best outcomes and allows you provide even better decision support to your knowledge workers in the form of a 'Best Next Action'. Best Next Action provides the suggested next action to the user but does not automate the decision. It is still up to the user to decide on the course of action.

At this stage, Predictive Analytics will bridge the gap between your data and actions by predicting both the asset failure events and the best actions to take in response.

SCALE PDM PROGRAM ACROSS MULTIPLE ASSET TYPES

In this step, you'll visually configure event streams in XMPRO for more of the assets you identified in Step 1. Unlike most predictive maintenance solutions which only allow for the monitoring of a specific type of asset, like a pump or turbine, XMPRO allows you to set up multiple monitoring queues for multiple types of assets. This makes it easy for you to scale your PdM program without having to code a bespoke solution for each asset type.

MOVE FROM PREDICTIVE TO PRESCRIPTIVE ANALYTICS

Evolve from a purely predictive to a more prescriptive program by adding business rules and parameters from other contextual data into your predictive maintenance event stream. An example of this would be:

If the asset is worth more than \$1 000 000 -> automatically escalate the case.

Prescriptive analytics is often referred to as the holy grail of analytics programs, which can make it seem like a completely unattainable feat. However, if you start by systematically implementing the steps outlined in this roadmap and strategically investing in your IoT-based PdM program, your organization will rise above the rest and see the benefits of the Internet of Things reflected in your bottomline.

“Never try to solve all the problems at once — make them line up for you one-by-one.”

— Richard Sloma

ABOUT XMPRO

XMPro is a global provider of software and services for industrial IoT applications.

Founded in 2009, XMPro has a proven track record of helping organizations improve operational excellence, reduce risk and increase asset utilization through its Agile IoT Application Suite and Business Process Management products. Closely integrated with the .NET and Azure™ software stacks, XMPro enhances the productivity of the Microsoft environment with a visual, model-based approach that enables deployment of IoT applications in weeks.

HISTORY

XMPro founder, Pieter van Schalkwyk, identified a need for an intelligent business process management solution in 2009. XMPro released its first BPM solution in 2010 and received the Gartner BPM Cool Vendor award in 2012.

Building on its BPM roots, XMPro evolved into an IoT company in 2015. It released the Agile IoT Application Suite to help industrial organizations improve operational performance through the Internet of Things. XMPro also joined the most progressive and innovative companies from around the world as a member of the Industrial Internet Consortium.

RECOGNITION



CONTACT

For more information visit xmpro.com or send an email to info@xmpro.com