SENSEYE PREDICTIVE MAINTENANCE

How To Run a **Predictive Maintenance Proof of Concept**

Best practice, common pitfalls, and how to ensure success.

Find out more: **siemens.com/senseye-predictive-maintenance**

SIEMENS

Introduction

AI and analytics are hot topics and almost every vendor claims an ability "to do" Predictive Maintenance with their predictive analytics platform. However, the relationship with a Predictive Maintenance software vendor is typically lengthy and requires a large amount of trust. With a large amount of noise in the market, it therefore makes sense to perform an exercise to down-select potential vendors. As with choosing your doctor, you need to be certain a vendor can do what they say they're capable of.

It often makes sense for an organization to run an exercise to ensure they not only select the right vendor, but that the organization itself is sufficiently prepared to ensure a Predictive Maintenance project has everything it needs to succeed, and that it is able to change and adapt to take advantage of the benefits on offer. The Predictive Maintenance proof of concept (PoC) can help to ensure this – if done the right way.

In our experience, many companies have tried and failed multiple Predictive Maintenance PoC's before they finally achieve the results they expect. Some of our customers have tried three other solutions before achieving what they expected. The lessons learned from the failures were that often the vendor wasn't entirely at fault – the problem actually lay within. It was only their awareness of success in other industry projects that gave them the confidence to make changes internally and persist to achieve success.

In this white paper, we try to share some of the lessons we've learned from both successful and unsuccessful PoCs. We'd love to be able to say that all our PoCs have achieved their end goals of scaling, but this would be ignoring valuable insights that we can now share to help you to avoid making the same mistakes.

Predictive analytics ≠ predictive maintenance

It's important that we address a key and common misunderstanding: predictive analytics tools can be used as part of a Predictive Maintenance program, but predictive analytics and Predictive Maintenance are far from the same thing.

Data scientists can work on machine data and identify anomalies and trends – and produce some convincing screenshots. But it's another thing entirely to have a deep understanding of what those things mean to the health of machinery, and to be able to have an in-depth discussion about what maintenance strategy to adopt with this new information. Pure data scientists usually lack the experience and knowledge needed to act as maintenance engineers. As a result, the custom algorithms they devise often perform poorly in real-world industrial conditions.

Poor performance in the real world: A good example of this is over-reliance on a NASA Turbofan dataset for demonstrating Predictive Maintenance prowess. It's an excellent dataset and can be used to demonstrate some very convincing prediction results. The problem is that the data is too good – it doesn't represent noisy real-world industrial conditions. Indeed, we've found that algorithms that perform well with this dataset perform poorly in the real world. We've stopped interviewing people who mention they have performed the bulk of their work using this dataset.

Key takeaway: Ensure the data you use is from your real world – warts and all. In other words, Predictive Maintenance solution vendors must be able to demonstrate that they can deliver in practice, not just in theory.

Spotting a vendor more used to predictive analytics than predictive maintenance

Special machines may have special and unique failure modes, but the failure modes and the kinds of information needed to detect them in most common machinery, such as motors, gearboxes, and robots are very well understood from a condition monitoring perspective. Failure often "is what it is" so, if a vendor is asking basic questions about machine failure modes and asking you to define everything, it's clear they don't have a background in condition monitoring or machine maintenance and don't know what they're doing. Chances of success are very slim.

Questions to check and clarify things are good. But questions such as "what does this vibration signal mean?" are a warning sign.

Asking you to label everything at an early stage – and being unable to understand what's shown in your maintenance logs and make correlations – can mean you're dealing with a vendor who's going to be taking a bespoke modelling approach. This can have very good results for up to tens of machines but will struggle with scalability and cost for anything more.

What is predictive maintenance?

Predictive Maintenance is the use of on-line monitoring to estimate the condition of machines. It is much more than applying condition monitoring in your maintenance strategy as you might be doing with Condition Based Maintenance (CBM). Importantly, Predictive Maintenance sources data during normal operations, hence minimizing disruption to operations for sampling or measuring. Traditional condition monitoring comes with a human burden on data collection and analysis. Predictive Maintenance removes this dependence, hence removing barriers for the application of cost effective condition monitoring to a much wider range of asset criticality.

Predictive Maintenance and CBM are complementary disciplines. If you already have a CBM program then culturally your organization understands the benefits of monitoring machines to optimize your maintenance policy. Predictive Maintenance takes that to the next level.

Predictive Maintenance utilizes a much higher level of data driven approaches than traditional condition monitoring as the data is from on-line sources opposed to irregular manual readings. The ability to be data driven enables Predictive Maintenance to exploit techniques form the domains or artificial intelligence and machine learning. Predictive Maintenance can deliver a range of game-changing business benefits including:

- Saving money by reducing the need for regular inspections and preventive maintenance
- Helping to plan corrective maintenance to minimize disruption
- Preventing unexpected failures
- Slashing unplanned downtime by up to 50%
- Extending the planned lifetime of assets
- Improving maintenance efficiency by 30%

Best practices

We've explained some key points for you to consider in order to maximize the chances your investment in this exercise gets you the results – and the vendor – you need.

1. Target specific business benefits

A genuine Predictive Maintenance solution must deliver – and be tied to – genuine business outcomes. An effective PoC must therefore be able to demonstrate how any proposed deployment will impact on business success.

Some vendors will have existing case studies that provide evidence to support their claims. Others need to rely solely on the results of a PoC to make a convincing case. In all of these, ask the vendor what the specific business outcomes of their projects were and ask to speak with existing customers who can corroborate these. Enough success studies are now in the market that a vendor who cannot produce a willing customer reference has likely only previously run 'science projects'.

2. Know your KPIs

Ensure you select a KPI that will move during the PoC. KPIs should improve, whether that means a reduction in downtime or higher overall equipment effectiveness (OEE), for example:

- Unplanned downtime reduction: This can be a big one and should be relatively easy to spot. There are some industries, however, where this isn't such an important factor and so other KPIs need to be chosen.
- Labor spend: The amount of labor devoted to condition monitoring should be reduced.
- OEE improvement: Not applicable to all, but an improvement in machine health will yield production efficiencies through improvements in performance, quality, and availability.
- Assets measured per person: The number of assets that can be managed by each reliability technician or engineer should increase substantially. With the right solution, it should be feasible to monitor thousands of assets per person and there shouldn't be significant effort involved in adding extra assets.
- Preventative maintenance reduction: The volume of preventative maintenance activities should start to reduce as maintenance teams rely increasingly on the Predictive Maintenance solution to direct their efforts more precisely.

3. Engage the users

It's not all about the vendor. Your users have the biggest role to play in the success of any Predictive Maintenance deployment. Trust and transparency are key. Vendors therefore need to understand what the user wants to achieve, so customers need to be clear about their goals in undertaking the project.

User engagement is a primary consideration when it comes to determining how useful a Predictive Maintenance solution will be in the longer term. Without user buy-in, even the most valuable tools will end up being shelved.

4. Make sure it's live – not smoke and mirrors

Some vendors expect prospective customers to accept a desktop demonstration as a PoC, but this sort of simulation can't possibly provide all the evidence needed to showcase the business benefits and user engagement a live deployment can demonstrate.

Some might argue that feeding in historical plant data and waiting to see what the system spits out can at least prove the underlying algorithms are robust. Even this could be in doubt, however, since there will be a temptation to "polish up"' the screen shots before sending back the results, leading to unrealistic expectations.

An initial deployment using live data, and that's usable by you and others in your organization with little to no training, is the gold standard for a PoC. It's the only way to demonstrate that the Predictive Maintenance solution can work "out of the box"' and successfully extract useful information from live plant data – however noisy it might be.

Ideally, live data will be used but this isn't always feasible. Sometimes, the best and representative results require up to 12 months' worth of historical raw data to be fed into the Predictive Maintenance solution to have a good chance of showing interesting changes and results.

5. Get the right data

Related to the above point, you will have heard the adage "rubbish in, rubbish out", and this absolutely applies to Predictive Maintenance given its reliance on online data. The data you share with the vendor needs to be in an unfiltered form from the assets – ideally in as raw a format as possible – and it needs to truly represent what's happening to those assets. Data collected at the wrong time and in the wrong way can throw everything off – if your internal condition monitoring experts can't see anything useful in the data then it's unlikely anyone else will be able to. Predictive maintenance analytics is not magic!

Variety is key: Several tens or hundreds of machines should be used. A truly representative sample with a variety of machine types will allow you to check that the solution can handle data from a range of asset classes without significant effort from you or the vendor. If they can't handle this, then they're likely designing bespoke models each time, and this won't be a scalable enterprise-wide approach.

Don't forget the context: Machine condition indicators can be interesting by themselves, but they're more meaningful with some context, such as maintenance events. Include these in order to help your solution provider with their setup and onboarding.

6. Experience check

Some vendors are too focused on the predictive analytics and do not have the industrial maintenance experience to design a useful Predictive Maintenance system. Potential customers can learn to spot the warning signs.

For example, it's fine for a vendor to ask a few questions for clarification, but if they're asking basic condition monitoring questions it suggests they may lack necessary plant maintenance experience. After all, a Predictive Maintenance deployment should be driven by maintenance and operational needs, not data science.

Many assets suffer from many of the same common maintenance headaches, such as worn bearings, failing seals, or misaligned moving parts. Customers shouldn't have to spend much time training vendors up to understand their machinery.

7. Scalability, scalability, scalability

Last but certainly not least, scalability is key. If the Predictive Maintenance system needs to be customized or tuned too much to suit each new asset or class, it will put the brakes on a rapid deployment and place a severe strain on engineering resources. An enterprise-wide roll-out is a huge undertaking, so you should also be asking how many thousands of assets per week the Predictive Maintenance vendor can onboard and how much effort that requires from you or the vendor to support and maintain. A Predictive Maintenance rollout with 10 machines is very different to one with 1,000, and there are very few suppliers who can manage that effectively without placing a tremendous burden on you and your colleagues.

Watch your back – **Common pitfalls on the customer side**

Much of this white paper discusses best practices with the Predictive Maintenance vendor but these are only half the story. It's important to highlight that YOU as the customer have an extremely important part to play in the success of a Predictive Maintenance project. This is not an exhaustive list but consider this helpful to drive further thoughts that in Senseye Predictive Maintenance's experience, all successful customers should ensure they avoid the following crucial points:

The wrong driving force: When the project isn't driven by a requirement from maintenance or operations, misunderstandings happen. Data science has a lot to contribute but let the experts in maintenance lead the maintenance project.

Ignoring scalability: By relying on manual models and tuning, this approach works well for a small number of assets or if you have extensive in-house data science expertise, but it can't be scaled economically or effectively.

Not considering the effort to scale: If you're looking to roll-out a system enterprisewide, how many thousands of assets per week can the Predictive Maintenance vendor onboard? A number in the tens per month just won't suffice.

Only focusing on critical assets: The supporting plant equipment has a very high chance of causing major unplanned downtime through "Swiss cheese" effects, and the maintenance effort on all of these ancillary systems for planned interventions and inspections is high. Consider the impact these assets have – it might well be more than just the original critical equipment you were considering.

Nobody wants to use it: User engagement is key – maintainers need to be excited about using the Predictive Maintenance solution and need to understand that it isn't a threat to their jobs and, instead, will only make them more productive. Internal cheer-leading should support this.

Lack of opportunity: Successful customers come with an expectation of what and where they can save if they can automate work they're often already doing. Make sure your scope is well-defined, as you don't want to change that mid-project.

Not having a goal: Maintenance isn't a science project. What are the KPIs you're looking to achieve? For example, reduction in unplanned downtime or lower maintenance spend (if so, by how much?).

Not being transparent with the vendor about your goals: They want you to achieve success and will move heaven and earth to help you do this. But how can they do this if you don't tell them what your goals are?

In summary

Maintenance is the ultimate practical 'hands-on' discipline, while data analytics is rooted in application of theory and advanced mathematics. The PoC should be the point at which those two worlds converge to deliver measurable business benefits through Predictive Maintenance.

Finding the right Predictive Maintenance supplier is tough – they need to understand your maintenance practices, your machines and what business outcomes you are looking to achieve and they need to be as committed as you are to achieving those. They need to do this through application of the right theories and mathematics – and do so in a way that is economically and organizationally scalable.

Senseye Predictive Maintenance: a risk-free solution

Senseye Predictive Maintenance is a cloud-based, software-as-a-service solution backed with expert consultancy and onboarding guidance to provide a whole"'turn-key" Predictive Maintenance package.

Existing users include blue chip companies in manufacturing, heavy industry, automotive and FMCG, who typically enjoy a 50% reduction in unplanned downtime and achieve a full return on investment in under six months.

By helping users to identify where they should be focusing their maintenance resources, Senseye Predictive Maintenance enables true Predictive Maintenance and delivers major productivity gains:

- 50% reduction in downtime
- 55% increased productivity
- 85% increase in maintenance accuracy
- Slash unplanned downtime by up to 50%

Published by
Siemens AG

Digital Industries
Customer Services
P.O. Box 31 80
91050 Erlangen, Germany

For the U.S. published by
Siemens Industry Inc.

100 Technology Drive
Alpharetta, GA 30005, United States

Article No. DICS-B10147-00-7600 PDF 0221 DÖ

Find out more: **siemens.com/senseye-predictive-maintenance**