E-BOOK

Data-driven Predictive Maintenance

The indispensable value of integrated data and advanced analytics to a comprehensive predictive maintenance and asset management strategy

By David Socha, Vice President, Industry & Business Consulting, APJ

06.21 / ANALYTICS / VANTAGE

teradata.

Table of Contents

Synopsis

This e-book discusses the critical importance of optimizing scalable, integrated data and advanced analytics as part of a credible and valuable predictive maintenance strategy.

The potential rewards of data-driven predictive maintenance are sustainable, available, and repeatedly proven. Data volume and availability regarding an organization's assets, and the entire business value chain, have never been greater—and continue to grow dramatically. But exploiting this data to deliver value requires significantly more than dashboards reporting on the operational status of assets, or even point analytics solutions capable of finding leading indicators of future failures using sensor data.

Instead, the full value of predictive maintenance can only be delivered through integration of data from many sources along with analytics—at scale. This must include data from operational technology (OT) systems such as historians, manufacturing execution systems (MES), and other control systems. Data must also be integrated from enterprise resource planning (ERP) and other IT systems, along with third-party data, data that's publicly available, and data securely shared by vendors and partners.

This e-book considers a successful predictive maintenance strategy from three perspectives:

1. Reliable predictions
2. Tactical interventions
3. Strategic interventions

All of these aspects rely on one thing: data. More specifically, opportunities to reliably predict future states, respond to those predictions, and find ways to make faults less likely in the future all rely on access to integrated and diverse data sets related to both the assets and the business they serve. The e-book will demonstrate that without diverse, integrated data, analytics are valueless. Yet with integrated and accessible data at scale, all potential value and benefits of predictive maintenance—and beyond into wider asset management—are unlocked.

Finally, an Appendix gives an introduction to the modern analytics platforms capable of supporting the integration and analytical techniques described in this e-book.

"The e-book will demonstrate that without diverse, integrated data, analytics are valueless. Yet with integrated and accessible data at scale, all potential value and benefits of predictive maintenance—and beyond into wider asset management—are unlocked.

Introduction: Predictive Maintenance in Context

Data-driven Asset Management

ISO 55000, the set of international standards for asset management and asset management systems, makes a clear and simple assertion on the purpose of assets. It states that assets exist to provide value to the organization and its stakeholders.

Not all assets are equal. In any asset-heavy industry, there may be assets of such high capital value they must be recorded on financial statements. In the same business, there will be assets of such low value they will almost be thought of as consumables. It is also highly likely that there will be assets and asset classes that are critical to have operational at all times and at (almost) any expense. Assets also include those that can be deployed then almost forgotten and allowed to run to failure; with many at some point in between. There will be assets that require frequent attention and others that require interventions once in a decade, if at all.

Modern asset management is wide-ranging, strategic, and data driven. We are in an era when data volumes and availability are growing at an unprecedented rate. For any enterprise reliant on physical assets, data about the capabilities, performance, value, and lifecycle of those assets touch every aspect of the value chain. This may be directly as a fundamental element of a department or division's success, or indirectly as a contributing factor.

Clearly, asset management cannot only be the responsibility of the asset management department. It is an enterprise-wide responsibility. Likewise, no credible asset management team will only consider what has traditionally been thought of as purely "asset data."

ISO Standards for Asset Management

ISO 55000 standards cover asset management. One example is in ISO 55002, Section 7.5.2, where a very broad range of information requirements are laid out, going considerably further than simply considering the efficient day-to-day maintenance, repair, and operation (MRO) of an asset base:

"In general, the organization should consider its asset information requirements related to the following areas

a. *strategy and planning (e.g. corporate service levels and objectives, asset strategy(ies), demand management strategy and plans);*
b. *process (e.g. process performance objectives and indicators, asset related processes and procedures);*
c. *technical and asset physical properties (e.g. asset attributes, ownership, design parameters, vendor information, physical location, condition, in service dates);*
d. *service delivery and operations (e.g. service levels, performance objectives, asset performance characteristics, future operational requirements, demand management objectives);*
e. *maintenance management (e.g. historical asset failures, betterment or replacement dates, future maintenance requirements);*
f. *performance management and reporting (e.g. asset performance data, continuous improvement objectives, regulatory reporting);*
g. *financial and resource management (e.g. historical cost, depreciation, asset replacement value, date of acquisition, materiality, capitalization rules, asset classification/hierarchies, life cycle costing analysis, useful lives of assets, residual value and any residual liabilities);*
h. *risk management;*
i. *contingency and continuity planning;*
j. *contract management (e.g. asset related contractual information, vendor information, service objectives, third party agreements)."*

Creating an Asset Management Strategy

Many business maintenance strategies have been developed without significant consideration of any strategic goals and without reference to any formal asset management strategy. Assets are typically maintained according to OEM specifications and guidance, or to fit with existing resource capacity and capabilities. This may produce acceptable levels of performance, perhaps assisted by high levels of redundancy and the general acceptance that "sometimes things fail." Even if this is the case, it is not an efficient strategy.

Modern maintenance strategies are risk and performance-based and are a fundamental pillar of a wider asset management approach like the one described in ISO 55000. Asset performance should be consistent with and directly contribute to the overall goals of the business. Investments to ensure those assets are capable of that level of performance should be calculated based on the value they contribute to the business goals, the cost of failure, and the level of risk the business is comfortable carrying.

To balance safety, risk, and cost, and ensure the efficient application of limited resources, the strategy must consider key elements such as the appropriate types of intervention for assets and asset classes. Should this asset be maintained, or left alone? Is time-based preventive maintenance the right approach for these assets? Should the organization, and could the organization, move certain assets to a predictive schedule?

Of course, these considerations are not a one-off exercise. Given any future change in circumstances, such as operational, technical, or availability of new data, does the overall balance of interventions remain correct or could it be improved? The strategy must also consider whether the interventions it details are at times inadvertently contributing to a reduction in performance. It is not unusual to discover that in some circumstances, maintenance, not operations, is contributing to early asset failure.

From Predictive Maintenance to Business Transformation

The fashionable concept of predictive maintenance has become almost synonymous with advanced maintenance strategies. This is typically a mistake—often perpetuated by software vendors, keen to sell their latest application or algorithm, or by sometimes ill-informed industry analysts promoting the latest, greatest, most-hyped use cases for the Internet of Things (IoT).

Asset management is clearly much more than a maintenance strategy. And an overall maintenance strategy cannot and should not be entirely based on predictive maintenance. However, there is considerable merit in exploring how data-driven predictive maintenance can be a key contributor to—or even the leading driver toward—a successful, comprehensive asset maintenance capability.

teradata.

Integrating and analyzing operational asset data with other data in the enterprise—and potentially from outside the enterprise too—can create insights that not only deliver reliable, actionable predictions for maintenance interventions but also inform every aspect of the business, from investment planning to risk management to human resources. As such, predictive maintenance can become the lighthouse program to drive the business transformation toward true enterprise-wide, data-driven asset management. Here's how:

A 3 Stage Approach: Predict, Respond, Prevent

At its highest level, the predictive maintenance concept is simple and familiar: if a company can predict when an asset is likely to fail, it can theoretically intervene before that failure. This leads to the potential benefit of increased asset uptime, which ranges from higher profitability to reduced risk of safety breaches. Even if a company is unable to intervene in time, it can at least plan for the event, make efforts to reduce any negative impacts, and ensure that normal service is resumed as soon as possible. Forewarned is forearmed, as the saying goes.

The potential rewards of predictive maintenance are real, available today, and have already been delivered in many businesses and industries. But achieving such rewards requires significantly more than simply being able to report on the operational status of assets or even finding leading indicators of future failures by using sensor data. In fact, a successful predictive maintenance strategy should be considered from three perspectives:

1. Reliable predictions
2. Tactical interventions
3. Strategic interventions

These perspectives could also be described as the ability to predict, respond, and prevent. Real success can only come from a combination of all three:

1. Can the company predict failure accurately enough to justify taking corrective action?
2. Does the company have all the resources needed to successfully respond to that prediction and take corrective action in time?
3. Does it make business sense to change certain practices and processes to prevent that kind of failure from occurring again—or at least make it less likely?

Thinking sequentially at first, it is clear that a reliable predictive model is a prerequisite for any justifiable action to maintain, upgrade, replace, or otherwise intervene in the operation of a working asset outside of what is dictated by OEM instructions and/or existing company procedures. Similarly, an understanding of the cost of failures and the value of successful interventions are prerequisites to any plans to change business processes or practices to make similar failures less likely in the future.

In reality, these are circular, iterative processes. The first prediction model does not need to use every possible data source or analytic technique. It must only be sufficiently accurate to demonstrate an improvement over what already exists. It may be so limited to only predict certain failures in certain conditions. But if this capability is valuable, it's already a successful model—as long as the business acts on the predictions and this delivers some value.

Prediction models can always be improved; tactical interventions can be further optimized, and strategic changes to how assets are maintained and managed can be reviewed and updated in light of new opportunities, locations, asset types, etc.

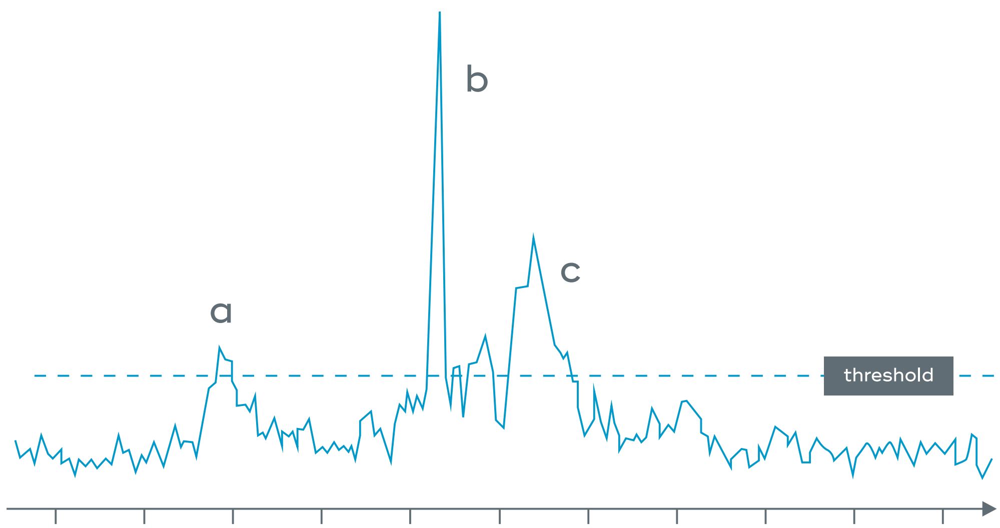

Figure 1. A simplified representation of sensor data

Predicting Failure and Taking Action

The Value of Sensor Data

Figure 1 shows a trace from a sensor on a rotating machine. It might indicate temperature, vibration, voltage, or any other relevant operational measure. It's possible that this trace predicts nothing at all, but it's also possible that with enough historical data in this format, the business could predict failures based on just the data from this single measure.

Pattern analysis could find that three threshold breaches in the style of (a-b-c) reliably predict the disruptive failure of the machine in three days, three hours, or three minutes. Similarly, analysis could find that the peak (b) is irrelevant and that the combination of peaks closely resembling (a) and (c) are reliable indicators of trouble ahead. Alternatively, if this is a trace over a period of months, perhaps each of the peaks could have been a failure in itself, subsequently repaired. In such a case, perhaps failure (c) always occurs when this class of machine has suffered failure (a) then (b) in a certain period.

It should be clear from this example that there can be significant value in the analysis of only operational data. Modern machine learning (ML) and other advanced analytics techniques can discover patterns, correlations, and indicators in the data that humans have failed to see. Models already exist that can predict certain failures and describe their likely root cause based only on correlations of sensor data.

ML will improve findings iteratively as the algorithms assess more data from the same source or by correlating it with new data from additional operational sources/sensors. Furthermore, ML models can be trained by subject matter experts to understand what is relevant and valuable, and what is, for example, a data quality issue as a sensor reports a reading that defies the laws of physics.

Applying advanced analytic techniques to operational data can certainly lead to good predictions of failures or degradations, and even show how particular failures may also themselves be indicators of different, future failures.

Figure 2 shows correlations between faults that cause shutdowns on a wind turbine. On the right, the analysis has been filtered to show how a shutdown caused by one particular parameter (Rotation Speed Unit Exceeds Limit) is related to other types of shutdowns and could be used as an indicator of future issues. This an example of faults predicting faults.

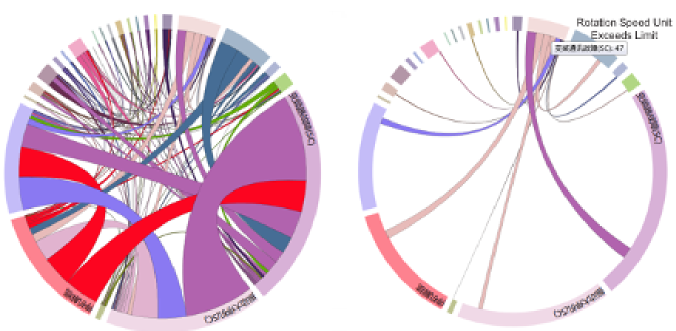

Figure 2. A fault shutdown causes additional shutdowns in wind turbines

Symptoms, Not Causes—The Limitations of OT Data Analysis

Given its potential, is the analysis of data from sensors, control systems, historians, etc., enough? In some limited circumstances, the answer could be yes. Analytical models using only operational technology (OT) data may prove quick and easy to deploy. They will be the responsibility of only one line-of-business department. They may even come as packaged applications, either standalone or bundled with other software from an OT vendor.

With enough warning to plan and execute some mitigations, these models may prove reliable enough to alert the business to future issues. They may be able to predict certain root causes of failure. And they may also expose initially inexplicable correlations that justify further investigation by engineers and data scientists.

Two things are certain. First, some predictions will only be accurate for timescales too short to plan anything other than a controlled shutdown. Second, OT-only models will only predict a subset of all types of failure or deterioration, and a subset of potential root causes. It is therefore critical to understand the fundamental limitation of this approach. With a few exceptions, failure prediction utilizing OT data focuses on symptoms of underlying issues in the asset base.

Overcurrent patterns, frequency fluctuations, excessive noise or vibration, temperature variations, etc., may be good indicators. But all are just symptoms of underlying problems. They are not root causes. Root causes themselves can also prove to be excellent indicators in an analytical model.

The Value of IT Data

We must recognize that both symptoms and potential causes of failure can be key indicators for a predictive model. In that context, consider a few examples of potential failure indicators not accessible via typical OT systems:

- Asset age
- Asset make, model, batch
- Maintenance, repair, refit, upgrade history
- History of which staff have interacted with the asset
- Geolocation and operating environment
- History of previous locations and moves
- Operating parameters vs. boilerplate specifications
- OEM or public data on performance in other businesses, locations, environments

Some of these examples may prove to have no significant effect on the operation of any asset in any location. Some may, like the OT parameters, also prove to be symptoms of failure. Others may prove to be root causes—perhaps a particular set of environmental factors in one or more locations contributes significantly to a specific recurring fault in a subset of pumps from one vendor.

Whatever the case, it is certain that once integrated with operational data, a subset of these will prove to deliver more accurate, more informative predictions of asset failure much earlier than looking at OT data in isolation. Consider a few possible predictions:

- Assets maintained by third-party contractors at two key sites fail on average 20% earlier than the same assets in other sites, and 10% more often than assets maintained by the direct workforce
- A particular specification of mobile pump from one vendor consistently fails one to two months after it has been moved from its original location to a new operating location
- One vendor's 50W DC motors' variation in Mean Time Between Failures (MTBF) has a high correlation with a specific range of manufactured dates
- The largest turbines from one vendor displayed performance deterioration up to one month before scheduled maintenance

Predictions like these have value at many levels—not only in creating more accurate models that cover more possible faults and root causes. Clearly, such predictions rely on analysis of data in traditional IT systems as well as OT. Businesses with a large fleet of assets may need to integrate and analyze very large and complex data sets. Companies can no longer rely on point-solution analytics. They need enterprise analytics at scale.

Only integrated, cross-functional IT/OT analysis can provide comprehensive prediction models that are able to reliably detect and predict asset failures arising from a broad range of contributing factors and—more often than not—make those predictions with lead times consistently long enough to plan and execute interventions.

Responding to Failure Predictions

Tactical Intervention

It should be well understood that for any business, there is little value in knowing an asset is going to fail if the business is incapable of intervening to prevent or mitigate the effects of that failure in the time available.

Operational systems should be able to detect potential failures in real-time and disconnect plant and equipment to minimize damage and maximize safety. But this should always be a last-resort measure.

Predicting failures with advanced analytics should ensure that more potential failures are mitigated long before protection systems are called upon. Consider three time windows for tactical intervention, based on predicted failures:

1. Immediate Shutdown

Predictions based on operational data may be able to alert operators to impending failures in time for them to safely shut down the assets before protection systems are called upon. This is not new. Such capabilities already exist in some control systems or associated OT analytics applications, for example. There are advantages to this capability. Switching out assets in a controlled manner is less likely to lead to asset damage or have safety implications. But critical processes may still be halted, service level agreements (SLAs) missed, significantly greater costs incurred than those directly associated with the repair, or replacement of the asset itself.

2. Planned Shutdown

The situation can be greatly improved if predictive analytics can provide sufficient lead time for operators to plan an outage. In many such cases, mitigations can be put in place to at least minimize the disruption caused by the asset coming out of service. In utility and telco networks for example, there are often opportunities to re-route the network to allow a particular switch or other item to be disconnected with no disruption to service at all. Even when full service cannot be maintained, a planned outage at least has the potential to schedule disruption at times when cost and risk can be minimized.

3. Scheduled Maintenance

The first two examples only consider how to take a potentially failing asset out of service safely and with the least disruption based on the time window made available by the predictive model. As explained, this has business value. But for greater value, companies must also be capable of safely returning the asset to service at the time that causes the least disruption. Often, although not always, this simply means as soon as is technically and safely possible.

Predictive models should therefore be able to predict failure with a long enough time window to allow companies to balance multiple factors such as risk appetite, resource availability, parts availability, and business/customer/partner/other stakeholder disruption. Only then will organizations begin to deliver on the real potential of predictive analytics in the asset world. In fact, managing these complex and interrelated factors should be considered the true beginning of a predictive maintenance strategy in any business.

Beyond Just Assets—Widening the Perspective

Delivering a successful predictive maintenance capability requires companies to enhance their models to predict failures with lead times long enough for the business to be able to carry out effective mitigations. Companies cannot rely only on operational data to do this. They must also find key indicators in data about asset age, location, repair, and maintenance history… anywhere that analytics uncovers data sets that enhance predictions.

But this is not the end. Companies must now realize that it's time to widen their perspective again.

Moving from an initial close focus on OT data to an understanding that also including asset-related data from IT systems is essential to create credible and comprehensive failure prediction analysis models will benefit companies. But they must now take the next step. From now on, they must widen their perspective again and see that improvements in other areas not directly related to asset data will be just as critical to the success of the nascent predictive maintenance capability.

To illustrate why, consider key factors involved in maintaining an asset that's predicted to fail:

- How long before the asset fails?
- Within that time, when is the ideal window for the outage to minimize disruption?
- Are engineers and technicians with the right skills and authorization available to carry out the work at this time?
- Are the right parts in stock?
- Can parts and labor get to the site in time?
- What is the cost of maintaining the asset before it fails?
- What is the risk and cost of the asset failing before this time?

Enhancing the predictive capabilities of the analytical model to extend the timescales of its predictions is certainly one way to improve a company's position regarding all these factors. But it is not the only one. In fact, it is highly likely that it's not always possible to reliably predict failures with the sufficient lead time to accommodate existing resourcing schedules or supply chain restrictions.

Answering Questions with Predictive Maintenance

A predictive maintenance program must address issues of the overall supply chain, logistics, resource availability, cost, and even risk appetite. For example, based on the best asset failure predictions, how can companies:

- Ensure critical components are in stock by shortening lead times or ordering earlier?
- Change the distribution of spare parts and consumables around distribution centers and local stores to bring them closer to where they will be required?
- Optimize shift patterns and locations of key repair and maintenance staff?
- Understand the cost of the proposed intervention versus the cost of allowing the asset to fail or resorting to a controlled shutdown?
- Understand the business' overall risk appetite to balance spend with acceptable risk and be able to justify the actions?

Delivering on the promise of predictive maintenance requires being able to answer these questions reliably and consistently. In such circumstances, integrated cross-functional data and analytics at scale are more than nice-to-have. They are fundamental.

Beyond Predictive Maintenance?

This e-book has discussed two critical components of a predictive maintenance strategy—the creation of reliable predictive models, and ways to ensure that tactical interventions responding to those predictions can be successful. It has covered how predictive maintenance and the wider maintenance and asset management strategies are inextricably linked.

But take a short detour from the simple 1-2-3 of predicting, responding, and preventing to consider what else organizations can do with a scalable, integrated, cross-functional data and analytics capability that's created to support a nascent predictive maintenance strategy.

Bottom-Up and Top-Down

In an ideal situation, a business reliant on assets will have credible, data-driven, enterprise-wide strategies for asset management and asset maintenance. Following a top-down approach, the overall business strategy will inform the asset management strategy, which in turn informs the maintenance strategy. In such cases, predictive maintenance will slot into this top-down regime as one method to deliver on those strategies.

The predictive maintenance program will encounter few barriers as it approaches other departments and divisions to agree on how to integrate data across traditional silos to deliver new insights into logistics, resources, and other critical components of the overall maintenance value chain. But many businesses do not find themselves in an ideal situation.

Many embark on initiatives such as predictive maintenance at a business unit, departmental, or Center of Excellence (CoE) level. This is a bottom-up approach without the benefit of top-down guidance or strategic support. It can only be successful for a limited time.

However, in such cases, the team responsible for delivering predictive maintenance has the ideal opportunity to lead the enterprise toward supporting its efforts by demonstrating—even facilitating—how business leaders can use integrated data and enterprise-wide analytics to create or enhance those top-down strategies that the entire business needs to succeed.

The very same data being integrated to enhance asset failure prediction, root cause analysis, resource optimization, logistics improvements, and the like is the data that underpins overall maintenance and asset management strategies. This is the data that will inform critical policy decisions on which assets should be moved to a predictive regime first and which later, which should remain on time-based maintenance, and even which make sense to run to failure—not because the company always has, but because the data says it is the right thing to do, according to cost models and risk appetite.

No matter what situation a business finds itself in regarding the status of overarching strategies and asset-specific initiatives, a collaborative, bottom-up and top-down approach based on data and analytics can benefit all parties. As such, a data-driven predictive maintenance strategy must recognize its part and influence in the wider field of asset management.

Preventing Failure

There is one more component to consider in the discussion of predictive maintenance—how to take strategic interventions that reduce or eliminate the risk of similar failures from happening in the future.

Consider a business that has deployed the analytical capabilities required to deliver value from reliable predictive models and has optimized key procedures to ensure it can perform the right interventions at the right time at the right cost. Figure 3 shows how some of the most important data sets may already have been integrated.

This hypothetical business has already integrated the data to deliver optimized tactical interventions. It is now time to take the next step and think strategically about what core business processes and practices could be amended to reduce or remove the risk of future asset failures. What additional strategic human resource factors could be influenced? What else could be done to improve the supply chain? And what about procurement? And site selection, location, and design? Investment planning? All of these have the potential to deliver a significant improvement in asset availability and performance if managed correctly.

> **“A data-driven predictive maintenance strategy must recognize its part and influence in the wider field of asset management.**

teradata.

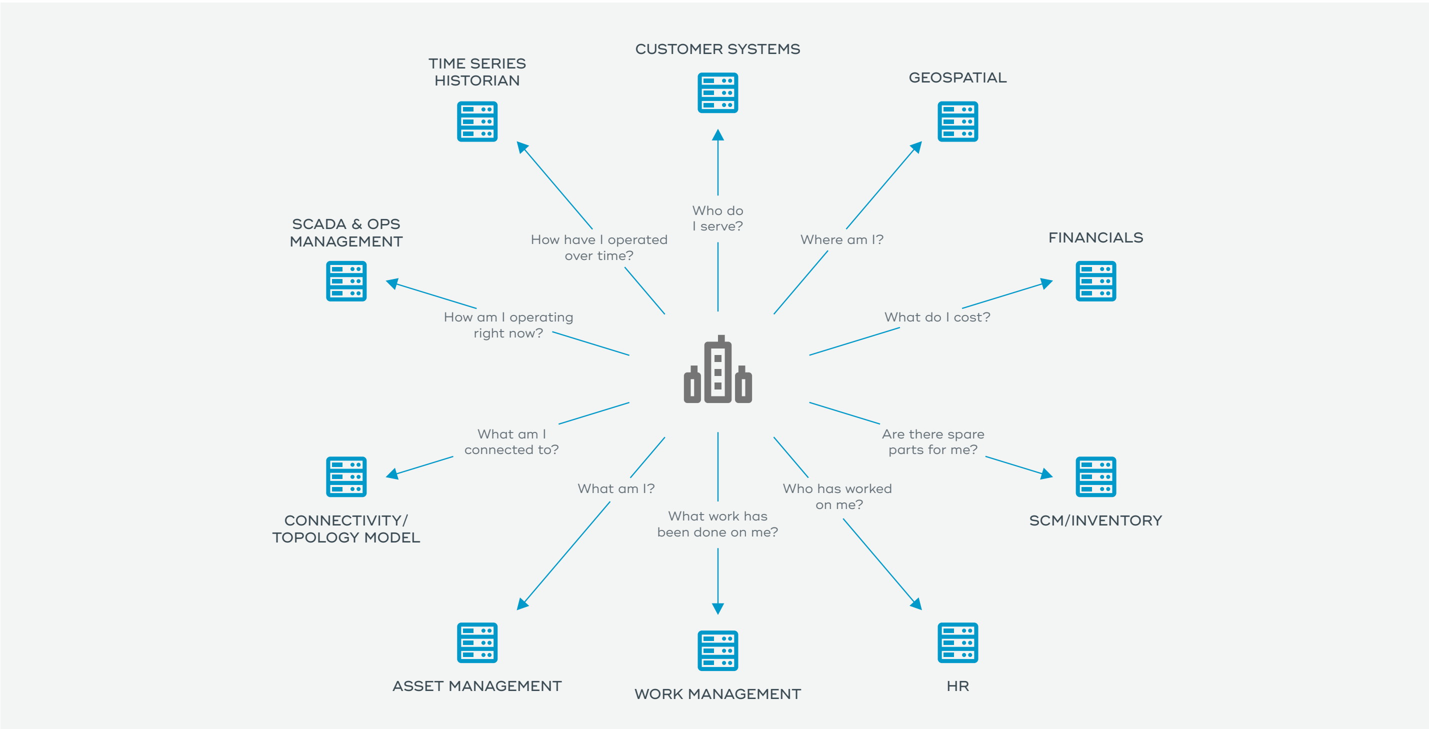

Figure 3. Indicative IT and OT systems associated with an asset. The diagram is not intended to be comprehensive. For example, it does not explicitly include weather or other environmental data. Data sources shown as separate may exist in modules of an integrated ERP.
The diagram does not include a reference to the business-critical data often found in individual or departmental spreadsheets, or in unsupported local servers in remote offices.

Examples of Strategic Interventions

This e-book's primary purpose is to explore the value of predictive maintenance. As businesses reach the end of their journey by predicting failures, responding to predictions, and arriving at a discussion of how they might prevent further failures from happening, they reach a logical (if blurred) boundary between any reasonable definition of predictive maintenance (Figure 4) and the wider topics of asset maintenance and asset management (Figure 5).

As such, the following section on strategic interventions is not intended to be comprehensive. Instead, it offers five example areas that can be reviewed, enhanced, and optimized using precisely the same integrated data and advanced analytic techniques deployed for predictive maintenance. The data and analytics can improve overall asset maintenance and asset management capabilities, strategically reducing the risk of asset failure. Many more strategic interventions exist.

1. Operating Environment

Similar assets operating in different environmental conditions may perform very differently. Mitigation options do not always exist. But in many cases, there are options to consider:

- **Indoor conditions:** Do assets in indoor locations perform better or worse than others due to local temperature, humidity, dust levels, or other controllable environmental factors? Would the benefits of improved asset performance in "bad" locations exceed the cost of improving the environment at these locations?
- **External shelter:** Do assets in external locations perform better or worse than those in enclosed or more sheltered locations? Would the benefits of providing shelter from rain, humidity, dust, etc., exceed the cost of the investment?

- **Asset relocation or replacement:** Are there assets from particular OEMs that fare better or worse than similar assets from other OEMs in certain environments? Is there a sound case for relocating existing hardier assets to more challenging environments, or bringing their less hardy counterparts to less challenging ones? Or is there a case for entirely replacing some assets not suited to their environment?

2. Operating Practices

Most often (and quite rightly) operational processes focus on safely optimizing yield, process efficiency, or other critical factors affecting the quality and value of the process outputs. Especially in the case of high-value assets, it can be worth considering whether such a focus is unnecessarily affecting the assets themselves:

- **MTBF:** Can it be demonstrated, for example, that reducing a production line's yield by a small percentage would still achieve targets and SLAs but would also reduce the burden on related assets, increasing MTBF, without further intervention?
- **Consumables:** Similarly, can it be demonstrated that reducing a production line's yield by a small percentage would still achieve targets and SLAs but would also reduce the wear on consumables, leading to fewer replacements and lower costs?
- **Asset lifetime:** Finally, can it be demonstrated that some similar small reduction to outputs that would still achieve targets and SLAs could increase the total expected lifetime of a high-value asset, giving the business an overall positive result?

3. Repair and Maintenance Staff

Analytics regarding how human resources interact with assets can yield surprising results and lead to many broad-ranging recommendations:

- **Correlating people and asset performance:** Are there particular crews that increase MTBF in assets they have maintained while other crews reduce it? Why is that? Do some crews work better with one asset class than another? Has the business identified a training issue? Or a lack of resources in one area leading to rushed work? Or another reason? How can the company bring everyone closer to the level of the best teams?
- **Resource planning:** Can the business see if it does not have enough skilled resources in one area (or in all) to carry out the work required to mitigate predicted failures? How does it address this issue to deliver maximum benefit quickly? Recruitment? Additional training of less-skilled staff? A short-term third-party contract?

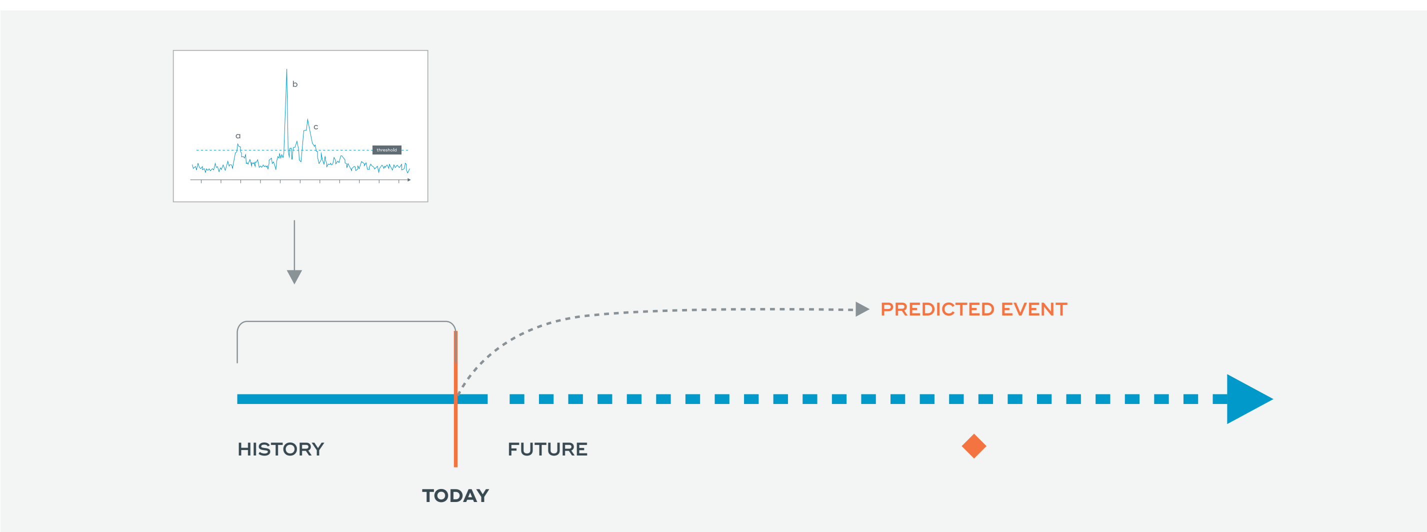

Figure 4. Predicted Event

- **Fraud detection:** Can the business see and stop suspicious patterns of activity? Are certain staff creating potential problems in the hope of lucrative overtime hours? Are contractors billing for work not completed, or for work that proves to be of suspiciously lower quality than direct labor or other contractors?

4. Supply Chain and Logistics Optimization

Tactical enhancements can reduce the negative effects of an inefficient supply chain on predictive maintenance. More strategic options include:

- **Inventory management:** Based on asset performance models, what stock is likely to be required, where and when?
- **High-value spares:** Is the cost vs. risk of holding high-value items the same as it was before? Can the business revisit decisions, reducing sunk costs without increasing risk?
- **Supplier partnerships and integration:** Given an increased understanding of stock and requirements, can the business strike better deals with suppliers? Can it also integrate with supplier systems, removing delays and potential human errors in ordering and fulfilment?

5. Procurement

In many businesses, new assets are purchased not by asset management, but by procurement. Its role is to find the lowest cost option meeting the specifications it has been given. With reliable asset performance data, that practice could be significantly improved:

- **TCO/TLV-based procurement:** Rather than relying on initial purchase price, can the business better understand which assets or OEMs consistently provide assets with the lowest total cost of ownership (TCO) or highest total lifetime value (TLV)? Can purchasing decisions be based on this information and justify those decisions to auditors and regulators?
- **Enhanced risk-based procurement:** Can the business add risk of asset failure statistics to its existing risk profiles for the OEMs its purchase from? How does this affect purchasing decisions? Again, does the business have the analysis to justify those decisions?
- **As-a-Service (aaS) models:** Does a new understanding of asset performance make the company better informed when evaluating aaS proposals from OEMs? Can it negotiate better rates? Or decide for or against a model based on analytics?

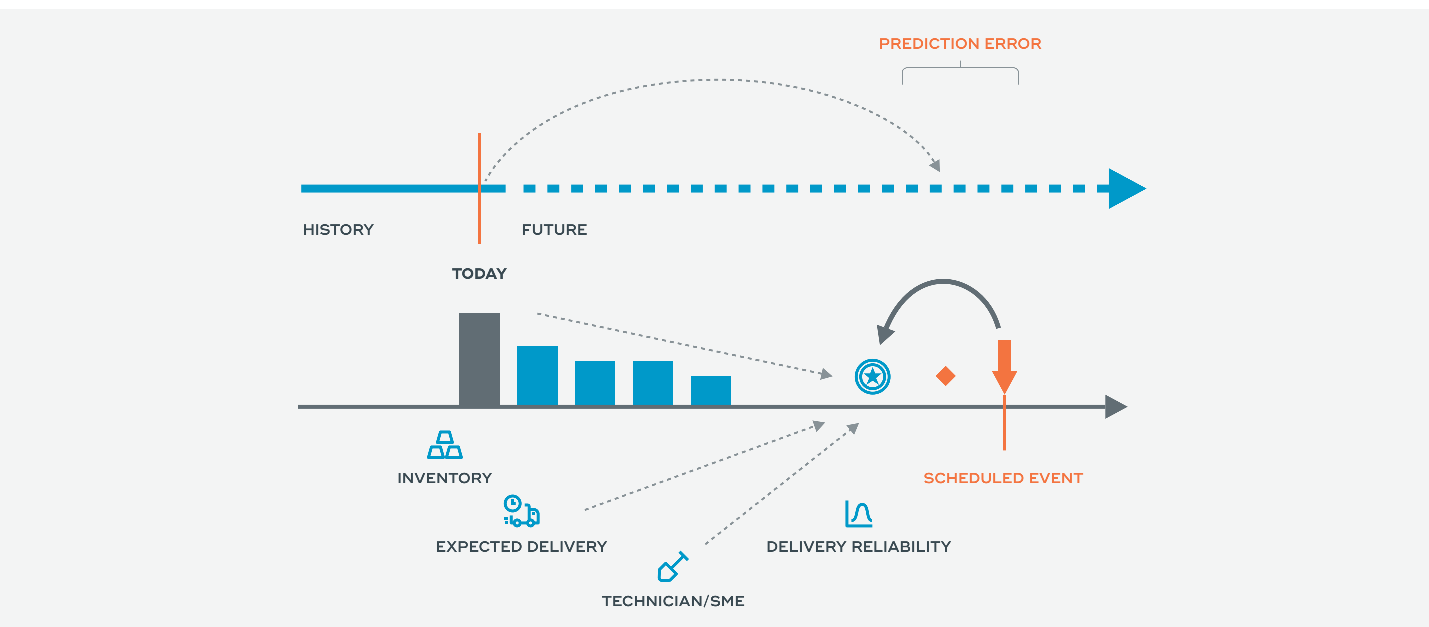

Figure 5. Linking Analytics to Actions

Gain the Full Benefits of Predictive Maintenance by Optimizing Data

Achieving the rewards of predictive maintenance requires significantly more than simply being able to report on the operational status of assets, or even finding leading indicators of future failures using sensor data.

Instead, the full value of predictive maintenance can only be unlocked through integration and analysis of data from many sources. The data comes from OT systems, of course, but also from ERP and other IT systems, and potentially even from third-parties, data that's publicly available, or data shared by vendors and partners. This requires a change in mindset from relying on departmental, point-solution, often application-based analytics to genuine enterprise analytics at scale.

As stated, a successful predictive maintenance strategy should be considered from three perspectives:

1. Reliable predictions: Can the business predict failure accurately enough to justify taking corrective action?
2. Tactical interventions: Does the business have all the resources to be able to successfully respond to that prediction and take relevant corrective action in time?
3. Strategic interventions: Does it make business sense to change certain practices and processes to prevent that kind of failure from occurring again—or at least make it less likely?

Advanced analytical techniques such as machine learning can and do provide answers to all of these requirements and more. Advanced analytics techniques already find the most important indicators of failure for different asset classes in a diverse array of industries around the world. They recommend the best courses of action and propose a schedule to minimize disruption.

Machine learning algorithms can learn how to balance a business's risk appetite, its budget limitations, the investment required to reduce asset failures by a chosen percentage, and the value that reduction in failures would bring. Without resorting to hyperbole, the possibilities are genuinely vast.

All of these opportunities rely on one thing—data. More specifically, these opportunities rely on access to integrated and diverse data sets related to the assets and the business. Without data, analytics are valueless. With integrated and accessible data, all of the potential value and benefits of predictive maintenance are unlocked.

teradata.

Appendix: Enterprise Analytics at Scale

The focus of this e-book is to highlight the fundamental role of integrated data and analytics at the scale required in large, asset-heavy industries with unprecedented and ever-increasing levels of data about their assets, people, partners, suppliers, environment… every element of their value chain. Its purpose is not to consider how data might be integrated or what analytics tools and capabilities are most appropriate to support predictive maintenance or other wide-reaching disciplines. However, it may be worth briefly introducing these topics for those businesses looking to begin a wider exploration of the subject.

Data Science, AI and Advanced Analytics

Many businesses place great faith in the discipline of data science or the promise of heady ideas such as artificial intelligence (AI). This undoubtedly has merit but is only part of the potential enabler for their ambitions. Data science and AI have their place. Both have the potential to be genuinely transformative, if applied correctly.

But both need the same critical, unavoidable foundation—integrated, trusted, diverse data sets at scale. Insights gained from AI, or by a team of data scientists working on sample data, can and will prove concepts, demonstrate new possibilities, find new correlations, and reveal undiscovered relationships that could lead to entirely new lines of business. The potential applications are endless. To realize actual business value, those concepts must be operationalized—at scale, in the real world.

It is informative to consider the following quote from Andrew Ng, founder of the Google Brain project, VP & Chief Scientist of Baidu, Co-Chairman and Co-Founder of Coursera, and an Adjunct Professor at Stanford University. In a recent lecture on deep learning (DL) and AI, he discussed the complexities of gathering data from many sources to enable DL and AI to deliver real value.

His advice was clear and simple:

"If your boss asks you, tell them that I said build a Unified Data Warehouse."[1]

The point is clear. DL, AI, data science, or any other analytic technique is only as valuable to businesses as the data that's available to analyze. Without data, there will be no insights. There will be no analytics-driven future.

The Analytics Platform is Key

As has been alluded to in this e-book, enterprise-wide analytics is a complex and multi-faceted subject, even when restricting the focus to asset-related analysis. A finance analyst working on better understanding TCO of ABB's pumps vs. Hitachi's should not be confronted with Google AI's TensorFlow for their analysis any more than a SCADA engineer will be able to analyze electrical fault data with business objects. Data scientists may wish to use a variety of the latest open-source tools, changing their preference as new business problems present themselves or new tools are released.

However, this cannot mean that each application or tool remains isolated, accessing its own silos of data, potentially duplicated many times across the enterprise. All users should be able to access the data they need from a single, secure, trustworthy platform that has the ability to scale to ensure performance at the enterprise level and serve users data in the format they require, via the tools that suit their role.

Figure 6 demonstrates the key elements of a modern, scalable, and extensible analytics platform. Its foundation is secure, scalable, and has highly accessible data storage optimized for analytics—one key element of Andrew Ng's recommendation to invest in a Unified Data Warehouse.

1 Nuts and bolts of applying Deep Learning (Andrew Ng): **https://youtu.be/F1ka6a13S9I**

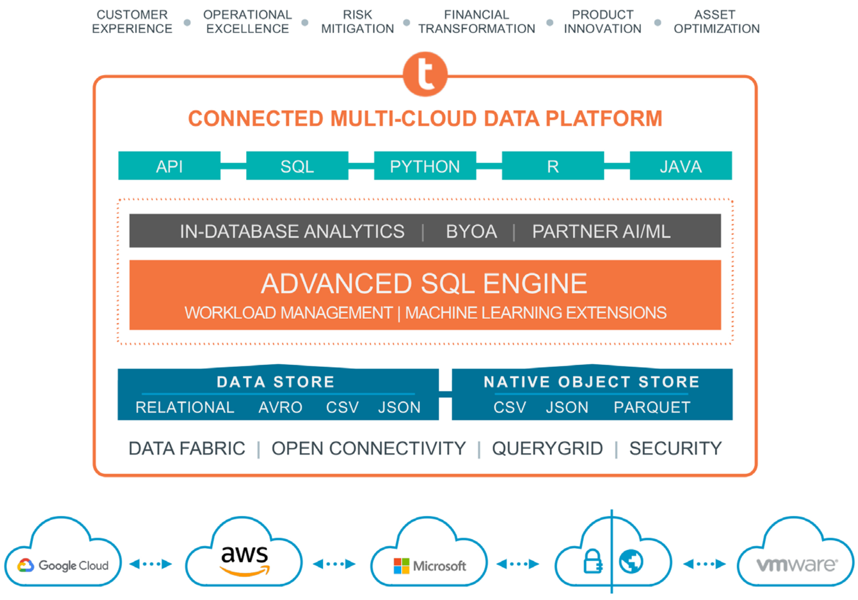

Figure 6. Key Components of an Enterprise Analytics Platform

Data must be accessible by many types of analytical engines, from the familiar SQL to the more modern and potentially esoteric, such as Tensorflow. Similarly, data scientists must be able to work with their favorite and most applicable languages and be able to apply new languages as they appear. Analysts familiar with SAS and similar analytical applications must be able to continue working with them as long as the business supports their use.

Beyond Figure 6, business intelligence users rightly unaware of the architecture they are accessing must still be able to run their reports on tools such as business objects. Line-of-business applications from vendors such as Siemens or GE must also be able to access this single source of truth, enriching the capabilities of their tools with the same secure, trusted, and accessible data used by the data science community and chief executives.

Such a platform is a major step forward in creating the data foundation that allows businesses to deliver the data-driven future they imagine today.

About Teradata

Teradata is the connected multi-cloud data platform company. Our enterprise analytics solve business challenges from start to scale. Only Teradata gives you the flexibility to handle the massive and mixed data workloads of the future, today. The Teradata Vantage architecture is cloud native, delivered as-a-service, and built on an open ecosystem. These design features make Vantage the ideal platform to optimize price performance in a multi-cloud environment.
Learn more at **Teradata.com.**

17095 Via Del Campo, San Diego, CA 92127 Teradata.com

teradata.