WHITE PAPER

On Behalf of

Making Predictive Maintenance Work

By Bob O'Donnell, TECHnalysis Research President

SUMMARY

To make an IoT-based project around Predictive Maintenance truly successful, not only is it important to understand the technical issues, but also the people and process requirements. In fact, many organizations with enthusiastic plans for Predictive Maintenance have found it difficult to move past the proof-of-concept stage and into full production because of challenges related to some of these issues. At the same time, organizations need to think about how they assemble their solutions, both technically and from a dataflow perspective. In this paper, we will discuss both sets of challenges and provide suggestions on how to overcome them, learn from best practices, and help ensure a successful deployment. Dell Technologies' comprehensive IoT portfolio cuts through the complexity and enables customers to access everything they need to deploy an optimized IoT solution. Working with partners enables them to provide complete and proven solution Blueprints to help customers reduce risks and get to ROI faster with IoT.

"The trick to a successful PdM deployment is to have a solid concept of what it is you want to achieve, and a well-constructed blueprint that can guide you towards the desired results."—Bob O'Donnell, Chief Analyst

INTRODUCTION

One of the most promising applications for the Internet of Things (IoT) is in the area of Predictive Maintenance (PdM). Businesses of all varieties use various types of machines in their day-to-day operations, and the idea behind PdM is to ensure that those machines—whether they're on a manufacturing line or in the building controls room — keep working. PdM using IoT is based on the premise that maintenance can be more timely, efficient and cost effective by "listening" to the machines themselves.

Of course, most companies have formalized maintenance programs for their critical equipment in place, but many of these programs are still very manual and time consuming. PdM is appealing because it has the potential to automate these processes and, by leveraging real-time analysis of sensor-based data, ensure that potential problems are discovered well before they actually occur. In addition, a well-defined PdM deployment allows companies to avoid the common practice of replacing machinery on a set schedule "just in case." Instead, companies can replace only the actual malfunctioning units, which translates into real-world monetary value. As a result, a predictive maintenance-focused IoT solution offers direct operational and cost-savings benefits to businesses, including reduced downtime, improved equipment effectiveness, increased return on assets, extended equipment life, and much more.

As enticing as PdM may sound, however, there are some potential challenges that companies who want to leverage the technology need to be aware of if they want to deploy it successfully. Some of these challenges are technical, some are operational and some, believe it or not, are actually organizational. The trick is to have a solid concept of what it is you want to achieve, and a well-constructed blueprint that can guide you towards the desired results.

ASSESSING NEEDS

The first step for any major technology initiative involves determining what exactly the end goal is and what steps or processes are necessary to achieve that goal. While that sounds straightforward in theory, in practice, this can prove to be challenging.

This is particularly true for IoT projects such as predictive maintenance. In addition to a primary objective (is this about saving money, reducing downtime, increased production or some combination of these and other factors?), consideration needs to be given to budget, timelines, security, processes and much more.

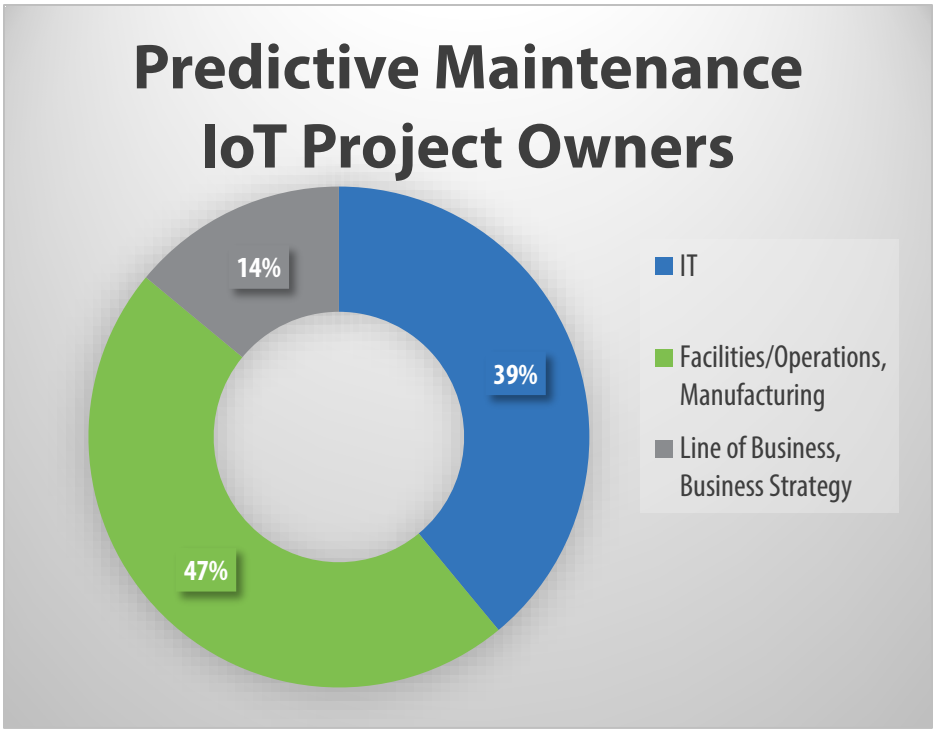

Another critical step is determining what parts of the business need to be involved. While most technology-based products tend to default to IT organizations, in the case of IoT, it's also critical to get operations groups involved. In fact, in most companies, more IoT projects will be driven by OT (including operations, facilities, manufacturing and more) than IT.

In fact, a recent survey of 620 IoT professionals (458 of whom had Predictive Maintenance projects) conducted by TECHnalysis Research showed that responsibility for Predictive Maintenance projects is split across different organizations in different companies, with classic facilities/operations and manufacturing groups owning 47% of projects, while IT owns 39%, and line of business professionals own the remaining 14%. Line of business groups can have a wide range of responsibilities, such as managing a division within a company's overall business, but like operations, they aren't typically the ones who purchase or manage large tech product deployments.

Operations vs. IT

These splits highlight what can be a somewhat sensitive subject in certain organizations: OT vs IT. Regardless of the quality of the relationship between these two groups, it's important to recognize that each of these organizations tends to work, operate and think differently about these kinds of projects. To be successful, and truly impactful on the overall business, it's important to acknowledge these differences and figure out a solid working plan.

Key to this exercise is determining which group is responsible for what because regardless of who ultimately owns the project, there's bound to be overlap and cooperation necessary in order to make a Predictive Maintenance project successful. For example, while OT may manage access to the machines being monitored, as well as determining the best way to acquire the right kind of data, IT may be responsible for providing the network connectivity and analytics. Ultimately, some group—perhaps a tiger team that incorporates members from OT and IT—needs to take ownership and responsibility for the effort in order to ensure that it gets successfully completed.

Scale

Another major challenge in most manufacturing environments is being able to obtain a consistent set of data across a range of different equipment and operating environments. Differences in the age of equipment, the manufacturer, the type of information the machine natively exports and many other factors can quickly turn what initially seems like a relatively simple idea into a very complex challenge. The process of scaling some initial concepts, in fact, can often prove to be a much bigger difficulty than coming up with the applications in the first place. As a result, it's essential to determine what the lowest common denominator in terms of data acquisition and analysis is—even if this makes the project appear to be somewhat simplistic.

Conversely, overeager consultants may start talking about completely automating an entire factory or production process, in the meantime losing what could be some relatively simple enhancements in the complexity of a bigger "solution." In practice, starting with a relatively simple goal and building from there is more likely to lead to a successful outcome. There's certainly no need to "rip and replace" in order to achieve a useful PdM deployment. Best practices suggest organizations start simply with what they have—perhaps one or two critical assets—and work on getting useful data from those machines before trying to expand the project.

For example, by deploying a gateway device, which offers onboard computing and connectivity options for machines and data networks, you can run software that essentially overlays some predictive models onto a factory floor without having to rip or replace anything.

Costs

As projects are developed, it's also important to consider financial issues, including concerns around the initial capital outlays. For some organizations, that might mean considering leasing options that turn capex into opex, while for others, direct purchase will be preferred. Interestingly, despite all the discussion around security issues with regard to IoT, the biggest

concern that IoT professionals who are involved with Predictive Maintenance projects have is the high initial costs for a project, according to the TECHnalysis Research survey mentioned earlier.

Nearly 60% of respondents listed this as a concern, vs. 45% for security-related issues. Part of this challenge is due to the fact that many projects take much longer to complete than initially expected, so organizations should plan accordingly. These results also highlight why it's important to partner with a vendor who can support "as-a-service" consumption models.

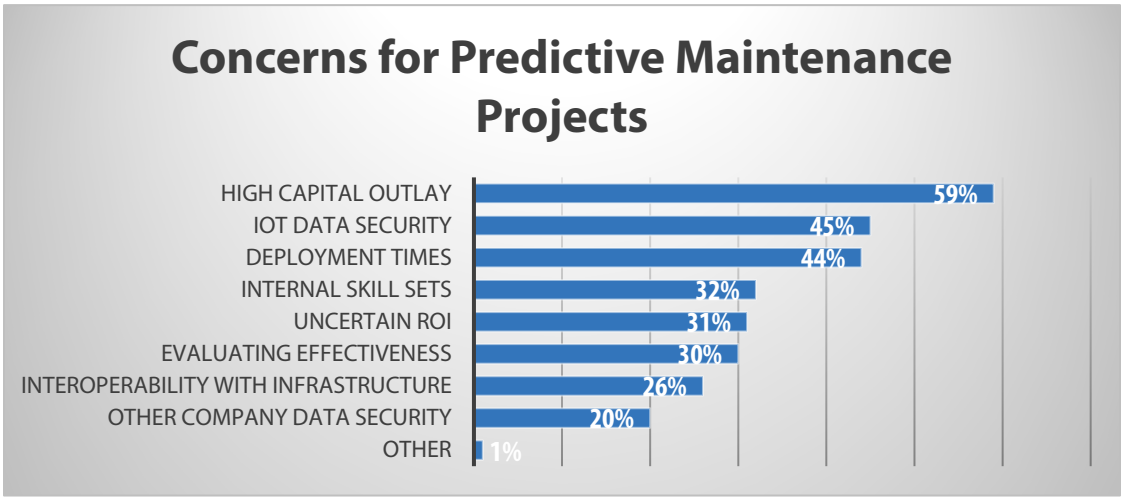

One key way to address several of these concerns is to think about building the right kind of computing architecture for a PdM deployment. For example, companies can potentially improve security and reduce costs by not trying to send all the data collected from all their machines back to a corporate data center. In many cases, on-the-spot analytics can be a more cost-effective and security-conscious solution.

To help make these kinds of decisions, it's important to have advice from organizations who've already been deploying these kinds of technologies. Dell Technologies has worked with many companies doing real-world PdM programs and condensed their learnings into a series of blueprints. These guidelines can help companies learn from best practices and make the right choices when it comes to their own situations.

Choosing the Right Architecture

From a technical solutions perspective, it's important to think backwards from the intended goal to the solution required to enable that goal. The key element there is determining the right type of architecture, particularly with regard to the elements needed, the connectivity available and the type of software expected to be used for analysis of the machine data.

One of the first steps in this process is figuring out what kind of data is already available and how you can access it. Then, it's time to ensure the network is there to bring that data into a computing environment. Finally, you need to determine where and how the computing and analytics should occur. In many instances, this is directly related to the type of machinery to be monitored and its native output. Dell Technologies works with partners like Kepware to provide software interfaces to newer equipment. For older equipment, companies can investigate add-on solutions that you can use to generate and then transmit sensor data from PLCs (programmable logic controllers), SCADA (Supervisory Control And DAta) systems and more to the computing environment.

Distributed Computing Model

In most situations, it makes sense to a build a distributed computing model, where some work is done on locally connected devices or gateways and some is sent either to corporate data centers or cloud-based services. These decisions will be determined, in part, by the type of existing infrastructure and/or business relationships organizations have in place. It's important to consider that while some IoT projects may be able to leverage an environment's existing computing devices, other projects may need their own or a different type of computing solution than companies currently use. The concept of "fog computing," for example, leverages the idea of using new gateways or small servers situated near the machines being monitored.

Another key benefit of a distributed model is that it allows organizations to create different levels of reporting depth that can correspond to organizational structures. For example, by having a computing gateway device like the Dell Edge Gateway 5000 situated near a key machine, you can generate and visualize basic data for the machine operator onsite. Simultaneously, you can send certain elements of that data up to a PC or small server in a factory floor manager's office, where it can be combined with data from other key machines to create a facility-level data dashboard with a wider perspective. Finally, that data and more from other factory locations or other sites can all be sent back up to the corporate data center or a cloud-based service where an even more complete dashboard can be analyzed by analytics experts.

For organizations looking to use machine learning or other AI (Artificial Intelligence)-based analytics tools, another benefit of a distributed computing model is the ability to transfer algorithm improvements from the data center or cloud back closer to the endpoint. Typically, the big number-crunching for machine learning leverages very large data sets and runs on large enterprise servers. This is called the "training" portion of machine learning. The end result is a series of pattern matching-driven algorithms that can be run on hardware devices, such as gateways, near the endpoint that respond to live sensor data—a process often called "inferencing." In many cases these algorithms are being updated on a very regular basis

through a feedback cycle in which new data is collected near at the gateway, funneled back to the data center, re-analyzed there, and then a new version of the algorithm is created and sent back down to the gateway and run there. Having a full-fledged distributing computing environment is what enables these powerful new applications.

Real-Time Data

Another critical element to consider when developing a plan for an IoT system's architecture is the type of data that is to be analyzed. In virtually all instances, some data points are more important than others and some are more time sensitive than others. As a result, it's important to determine what data needs to be acted on quickly—the real-time or "perishable" data—and what can be analyzed at a later time. In some cases, it may be determined that certain data points aren't worth capturing and analyzing at all, particularly in instances where enormous amounts of data are being generated.

A critical part of this process involves where and how data is stored and types of tools are used to analyze it. The combined storage offerings from Dell and EMC let organizations choose what they need based on their own unique requirements, from flash-based storage devices for speedy access, to larger capacity traditional spinning hard drive-based choices for deep storage and longer-term analysis.

On the software side, Dell partners with companies like Software AG and SAP to enable companies to look at both the structured and unstructured data that PdM applications might generate. From real-time sensor feeds to reliability models and maintenance logs, all the different data elements that are available can be leveraged to create solutions which find and prevent problems in critical machines before they occur.

Conclusions and Recommendations

The promise of the IoT has become so grandiose that some have jokingly likened it to world peace—it's something everybody wants and knows is a good goal, but getting there isn't quite so simple. Thankfully, achieving success with Predictive Maintenance-focused IoT deployments is much easier than world peace, but there are important issues that need to be addressed and dealt with along the way.

The key is to have a plan—a blueprint—that provides a logical path for your organization to follow as it tackles this challenge. Dell Technologies has a great deal of experience in working with companies of all types and sizes in creating successful PdM deployments and they've turned the knowledge they've acquired from doing those projects into exactly that—blueprints with validated solutions that can help your organizations get to ROI faster. By reducing the upfront time your organization needs to assess partners and stitch together

applications, these blueprints can help bring the clear benefits of PdM as quickly and efficiently as possible.

The other key point to remember is that it's not just about technology. One of the key differences with the IoT is that it forces several different elements of a business to work together, sometimes in a manner to which they aren't accustomed. Being cognizant of these potential organizational challenges is just as important as understanding the technological requirements, so be sure to consider them as you develop your PdM plans.

The real beauty of the IoT is that it can provide new kinds of insights into equipment, processes and procedures that might even be decades old. By smartly leveraging this insight, companies can turn any necessary investments into benefits that not only help the bottom line, but can improve efficiency and overall operational effectiveness. Done well, it's a real win-win.