The Essential Guide to

Predictive Maintenance

Applying Splunk and machine learning to your maintenance operations

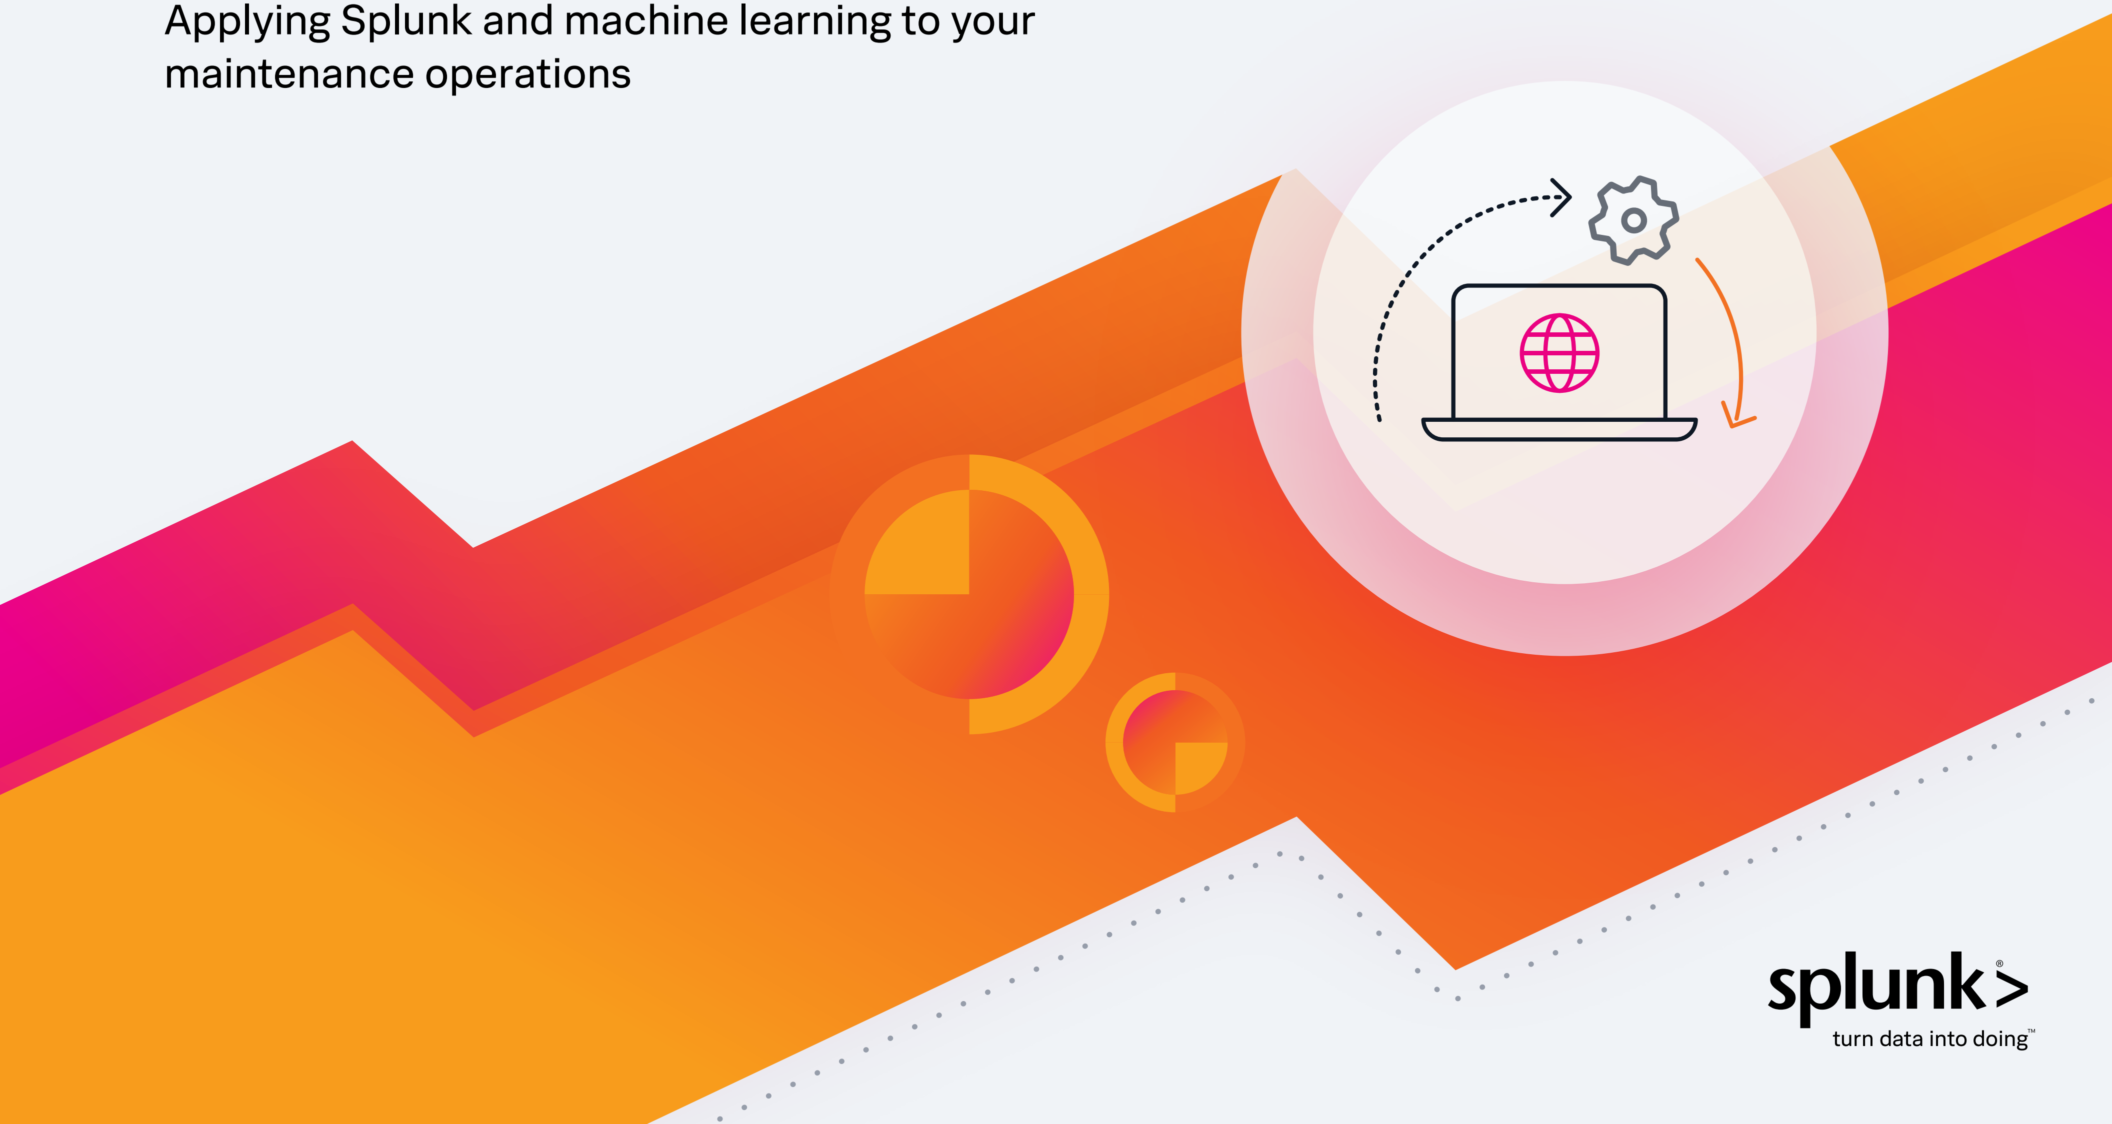

splunk>
turn data into doing™

Compete and Save With Predictive Maintenance

Table of Contents

Want to Know How Predictive Capabilities Will Change Maintenance?

With the ongoing explosion of inexpensive sensors and accessible big data technology, the way industrial organizations look at assets is completely changing. Critical, forward-leaning analysis of operations is no longer a matter of “what if” but of what can we learn, when and how fast. This is a boon for industrial operations, where opportunities that leverage operational technology (OT) data have the potential to disrupt overall bottom lines in just about every industry.

This is especially true of asset performance, as organizations shift from preventive to predictive maintenance. What company doesn’t want to be able to analyze data and predict when it’s the optimal time to do maintenance or predict failures of critical operational assets? It can lower costs, boost profits and improve the overall efficiency of the business.

By following Splunk’s four-stage journey, you’ll be doing predictive maintenance in no time.

Let's Break It Down

The ultimate goal for predictive maintenance is to optimize when and how often maintenance is performed on industrial assets with data. Armed with the ability to discover patterns and signals from sensor data, organizations can look around corners, apply timely maintenance strategies, and ultimately predict and plan for big future issues that cause massive downtime (and massive maintenance costs).

Planned maintenance vs. unplanned maintenance

Most organizations are concerned with two types of maintenance — planned and unplanned.

Planned or scheduled maintenance happens before a failure occurs — or prevents failures from happening. Unplanned maintenance, also known as reactive maintenance, happens when parts fail and need replacing. This type of maintenance is generally costly, as it affects operations due to asset availability, kinks in the supply chain and more depending on your industry. This also often means business grinds to a halt or slows down — and services are interrupted for customers.

Of the **82%** of companies that have experienced unplanned downtime over the past three years, those outages lasted an **average of four hours** and cost an average of **$2 million**.[1]

Unplanned downtime results in loss of customer trust and productivity — **46%** couldn't deliver services to customers, **37%** lost production time on a critical asset, and **29%** were totally unable to service or support specific equipment or assets.[2]

Forward-looking organizations are now seeking to more accurately predict problems instead of simply preventing them with standardized maintenance. It's about time too: 70% of companies are unaware of when equipment assets are due for maintenance or upgrade.[3]

Shifting from preventive to predictive maintenance

Preventive maintenance is planned maintenance to preempt failures caused by wear and tear. It's often expensive and not always necessary, although it is currently the safest option. The eventual goal — predictive maintenance — is also planned maintenance, but it takes into account estimated service intervals as well as data-driven insights based on the measurement of operating conditions, to monitor and diagnose equipment issues in real time. It's meant to catch anomalies in operations before they become major challenges that could affect the business. Predictive maintenance also makes the most of asset resources and extends an asset's life as proper and timely maintenance becomes more feasible, and eventually, integral to its life cycle. Still, it's no easy feat. According to a recent report, 24% of operators describe their maintenance approach as a "predictive" one, based on data and analytics. The rest either took a reactive or time-based approach.[4]

1. http://lp.servicemax.com/Vanson-Bourne-Whitepaper-Unplanned-Downtime-LP.html?utm_source=blog&utm_campaign=vansonbourne2017
2. http://lp.servicemax.com/Vanson-Bourne-Whitepaper-Unplanned-Downtime-LP.html?utm_source=blog&utm_campaign=vansonbourne2017
3. https://www.geoilandgas.com/sites/geog.dev.local/files/ge_offshore_study_paper.pdf
4. https://www.geoilandgas.com/sites/geog.dev.local/files/ge_offshore_study_paper.pdf

The Solution Defined

Industrial organizations, from manufacturing to oil & gas, need to be able to ascertain the state of their assets before it's too late. Companies need systems that conduct advanced analytics and machine learning so they can have the insights they need to predict issues instead of continuing to react to them. Fortunately, the ecosystem has finally reached the point at which it no longer has to be reactive or preventative. Condition-based and predictive maintenance are starting to change the industrial world.

The Splunk platform makes this possible. With all your data in Splunk, you can easily leverage Splunk machine learning tools to understand your data and construct machine learning models to start taking predictive action in no time.

Splunk for Industrial IoT simplifies data preprocessing through a series of exploration techniques to provide a clear picture of your assets, what the metrics mean in the context of the assets and their components, and how the metrics change over time. Armed with the appropriate techniques for your dataset and the insights they provide, you'll avoid costly repairs while maximizing the use and availability of the equipment.

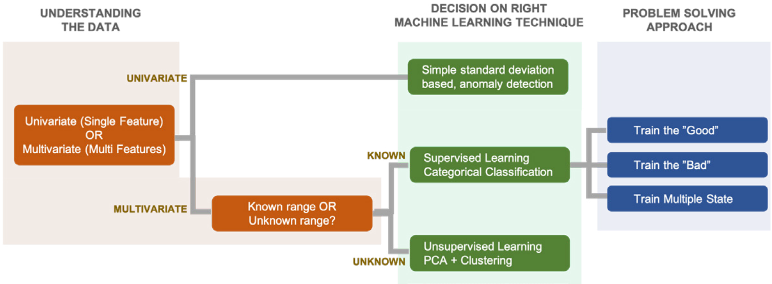

Want to See How You Can Get Started?

Let's walk through the process of effective predictive maintenance. We will also explore top use cases, helping you dive into what predictive maintenance can do in different settings so you can reach optimal performance.

Splunk's Stages to Predictive Maintenance

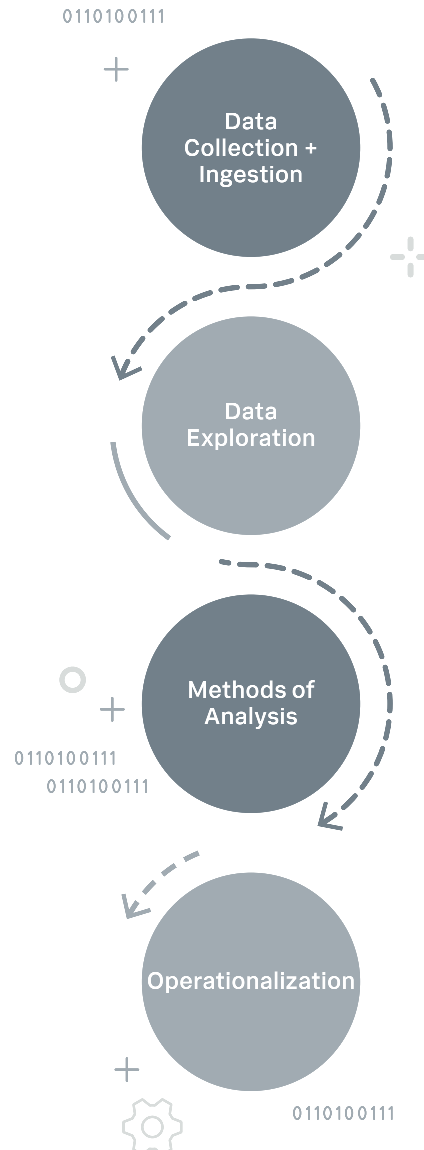

Stage 1:

Use Splunk software to collect, store and structure asset metrics

Stage 2:

Pre-processing and data exploration help you understand the type of data in use and the characteristics

Stage 3:

Conduct anomaly detection via supervised and unsupervised machine learning for predictive maintenance analysis

Stage 4:

Apply predictive models to broader implementations — create reports and alerts for operational actions

Next, we will provide you with specific industrial IoT maintenance use cases with which you can practice each stage. The goal: helping you understand and effectively navigate each stage in differing use cases so you can effectively make it a practice in your organization.

3 Major Types of Predictive Maintenance

Initial product defect and recall analysis

Summary: Not all products are created equal — whether it's vehicles, turbines, machinery or mobile devices — especially when they are first deployed. It's important to understand potential points of failure as these assets make it out onto the field or factory floor. With connected technology, this has now become feasible and almost necessary.

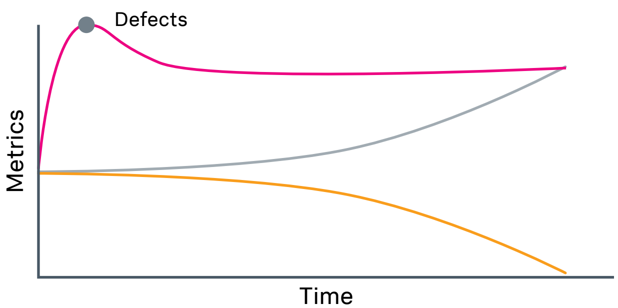

Unplanned maintenance by early detection

Summary: Downtime can be costly, and if operators miss the signs of failure, it can leave them scrambling to fix the problems. Early detection is critical so operations don't grind to a halt.

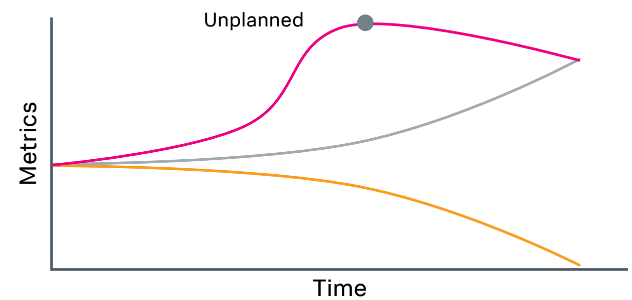

Recurring maintenance optimization

Summary: Preventative planned maintenance doesn't always mean everything runs efficiently. In fact, it could be the opposite if you're using resources on maintenance for assets that don't need it. Optimizing maintenance cycles through a data analytics approach ensures you do maintenance at the right time to maximize value and reduce cost.

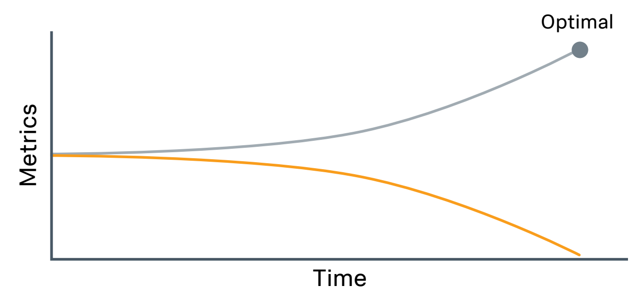

Getting Specific

Let's dive deeper into these use cases to get a better grasp of their importance and how you can leverage them in your operations.

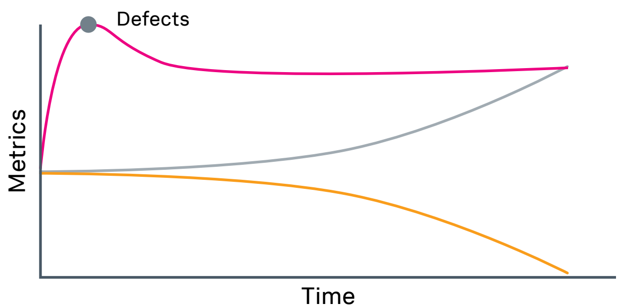

Use Case 1: Initial Product Defect and Recall Analysis

Description
In the example of heavy industry vehicles or any connected cars, monitoring for potential defects during operations can create opportunities for the manufacturer to not only monitor for any defects but also to personalize maintenance support services. All of this reduces vehicle failures and optimizes operations.

Why is it important?
Imagine an automobile manufacturer launches a new car model but soon finds out there is a potential batch of vehicles that could be affected by critical engine defects. With data collected from millions of cars, the manufacturer can analyze whether the defect exists or not. If it does exist, the manufacturer knows exactly which vehicles need to be recalled. This kind of insight can save millions of warranty costs.

How to implement?
Collect any performance metrics related to early life stage — typical metrics are temperature, pressure, flow rates or vibrations. The goal is to feature engineer the initial performance metrics that represent initial expected performance, then classify the defects against what's normal.

Methodology summary:
- Conduct analysis for an initial asset performance profile and understand what the expected initial performance metrics combinations look like.
- Engineer feature characteristics of an initial performance status from multiple assets.
- Some of the engineered features could be average, median, ranges and quartile analysis of multiple asset samples.
- Identify the initial asset performance that's positioned outside of major standard deviation bell curves.

One example may be collecting data from a vehicle's onboard mobile data terminal (MDT) and analyzing it via cluster analysis to group certain behavior across a fleet. Externally added sensors in commercial/construction vehicle fleets can also collect data vibrations, temperature readings and more.

Data sources
- Vehicle mobile data terminals provide relevant metrics.
- External sensors tracking variables like vibrations also provide relevant data.

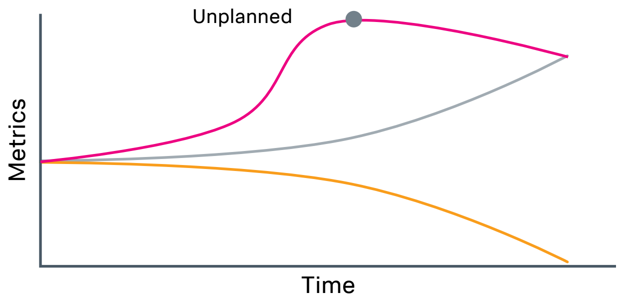

Use Case 2: Unplanned Maintenance and Early Detection

Description

Detecting downtime or equipment failures is another way machine learning techniques can be applied to get results. This also applies to monitoring asset performance metrics at a more granular level. Even if assets seem to have more life left, they can unexpectedly stop when a component fails. To detect unexpected failures, the analysis looks for signs and patterns that fall outside the norm.

Why is it important?

A large percentage of maintenance is related to unplanned downtime. It's costly, can affect operations downstream, and can damage a company's reputation with customers and users.

How to implement?

Collect any performance metrics related to asset life stage as well as typical metrics like temperature, pressure, flow rates and vibrations. The goal is to apply deviation analysis on metrics that change more rapidly with bigger dispersion so as to note any fast failing status for an asset.

Methodology summary:

- Conduct analysis for an asset performance change profile and better understand the changes in performance metrics over time.
- Calculate the R-squared (R2) between the expected increase rate and the actual increase rate.
- Identify the spikes in R2 between expected and actual events to define the threshold.

Data sources

Assets like turbines, motors, engines, compressors and more provide metrics for analysis.

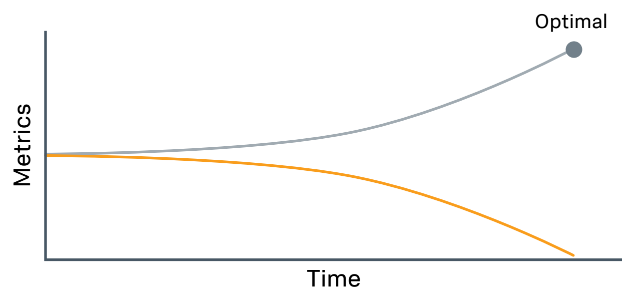

Use Case 3: Recurring Maintenance Optimization

Description
A pimary directive of an operations manager is to make sure assets are running at peak performance, while keeping costs under control. This means maintenance done right. Too much maintenance can be costly and too little maintenance could mean disaster. Optimizing maintenance cycles is key in getting this process right.

Why is it important?
Imagine you're an operations manager for a massive airline. Keeping your fleet's jet engines running is a key component to ensuring the availability and safety of your service. Recurring maintenance optimization plays an integral role in this process as failing to do it properly could mean up to $1 million a day in lost profits per jet.

How to implement?
Collect data from your fleet's engines while considering other factors like flight distances, altitudes and directions. The goal is to apply a data-driven, analytic approach to extend the maintenance life-cycle policy that would optimize the policy per engine, not for the fleet as a whole.

Methodology summary:

- Conduct analysis for an asset performance profile and understand performance trends/patterns of metrics as the assets wear out.
- Apply either anomaly detection, unsupervised learning or supervised learning to find the pattern of performance metrics that deviates from the norm or expected ranges of operation.
  - Anomaly detection approach:
    Detect outliers by applying a standard deviation method
  - Unsupervised learning approach:
    Cluster similar behavior performance metrics to group different conditions of assets
  - Supervised learning approach:
    Design a training model by labeling training data, and then applying the model to future date-sets to be predicted

Data sources examples
Jet engine metrics usually consist of 13 major sensors metrics including various speeds, temperatures, pressures and liquid flow ratios.

About **Splunk**

Splunk Inc. makes data accessible, usable and valuable to everyone. Download the Splunk Essentials for Predictive Maintenance and get hands-on experience with machine learning for predictive maintenance.

20-20901-SPLK-Essential-Guide-Predictive-Maintenance-104