



The Complete Guide to Industrial Maintenance Optimization

Table of Contents

Summary	3
Introduction	3
Predictive Maintenance	4
Imperfect Predictions – The Case for Prescriptive Analytics and Optimization	4
What is Prescriptive Analytics?	5
Rule-Based Decisioning Systems vs Optimization-Based Systems	5
Two Examples: Subway Car Maintenance Planning & Technician Dispatching	5
What is Optimization and How Does It Work?	7
Optimization Applicability Across Maintenance Planning Functions	8
Maintenance Optimization Use Cases in Industry	8
Heavy Equipment with Intermittent Usage - Aircraft Engine Lease and Maintenance Cost Optimization	9
Heavy Equipment with Continuous Usage - Aircraft Heavy Maintenance Optimization	11
Light Equipment with Intermittent Usage - Dynamic Tactical Elevator Maintenance Planning	12
Light Equipment with Intermittent Usage - Facility Maintenance Technician Dispatch and Scheduling	13
Conclusion	15

Summary

Maintenance optimization is a critical aspect of asset management, but it can be a daunting task for many organizations. With competing objectives of minimizing maintenance costs and maximizing asset uptime, it's not always clear how often maintenance should be performed or what factors to consider when determining the optimal maintenance schedule.

Fortunately, advancements in IoT, telematics, and sensor technology make it possible to employ new maintenance strategies. Specifically they enable a shift to more demand-driven and outcome-based maintenance planning. By leveraging these hardware technologies together with data-driven decision support systems powered by machine learning and optimization methods, organizations can better anticipate failures and ensure that maintenance activities are carried out efficiently and in alignment with organizational goals.

Introduction

Effective maintenance planning is essential for extending asset lifespan and maximizing uptime, but it's not always easy to strike the right balance between minimizing maintenance costs and minimizing the risk of business interruption due to asset failures.

While manufacturers typically provide maintenance guidelines, it's important to recognize that these guidelines may not be sufficient for your specific circumstances. Manufacturers may not have the same interest in maximizing asset life as you do and there are often factors beyond the manufacturer's control that can impact the asset's lifespan. For instance, cars located in coastal regions may require more frequent brake maintenance due to the corrosive effects of salt air and sea spray.

To optimize maintenance planning, it's important to consider a range of factors beyond just the manufacturer's guidelines. In the automotive parts example, this might include environmental factors, such as temperature, humidity, and proximity to saltwater, as well as usage patterns, such as city miles driven versus long-haul mileage.

In the following sections of this whitepaper, we will explore the key strategies and techniques for optimizing maintenance planning, including predictive maintenance, maintenance planning and scheduling optimization. We will also cover the latest AI software tools and methods powering the most advanced software systems that are used to support these use cases.

By taking a comprehensive approach to maintenance optimization and leveraging the latest technologies and best practices, organizations can reduce costs, minimize the risk of downtime, and maximize their return on assets over the long term.

Predictive Maintenance

One of the most common methods employed by analytically-oriented companies that maintain industrial equipment assets is to adopt a predictive maintenance approach. This involves using a variety of data, tools, technologies, and methods to predict when asset failures are likely to occur, allowing for maintenance to be performed proactively.

Modern sensors, IoT, and telematics devices can now provide real-time data on asset utilization or performance, which can be used to detect when maintenance is needed. For example, an electrocardiogram (ECG) machine can detect electrode degradation and report on this, enabling hospital technicians to replace the part before it fails.

In cases where assets don't have the ability to easily provide current utilization or asset status data, other data sources can be employed to estimate asset health or utilization. For instance, it may not always be feasible to directly track the utilization of elevators, automatic doors or food service equipment situated in a train station. A workaround could be to analyze passenger rail ticket sales by day, which can be utilized to estimate the equipment's utilization in the station, providing a more precise understanding of when maintenance should be performed.

To take full advantage of all available data sources, machine learning software methods are now commonly used to build predictive models that analyze all possible correlated data that could give insight into an asset's usage or likelihood to fail. These mathematical models can then be regularly run using the latest situational data and even demand forecast data to predict when asset failures are likely to occur, thereby informing the systems that decide when preventative maintenance should be performed.

Imperfect Predictions – The Case for Prescriptive Analytics and Optimization

The problem with predictive maintenance is that it's impossible to perfectly predict the future. Even if we integrate the most sophisticated sensors, external data sourced and predictive models, we eventually reach a point of diminishing returns in which the increase in prediction quality becomes marginal. This is why predictive maintenance alone is not enough. We need a system that can work with imperfect forecasts, create robust plans that incorporate some buffer for surprises, and allow for real-time adjustments as needed.

What is Prescriptive Analytics?

Prescriptive analytics goes beyond predicting the future. It is about using historical, current/situational, and predictive data and then using that to make specific decisions. These decisions can be based on everything from simple if/then rule logic to more sophisticated plans or activity schedules.

Rule-Based Decisioning Systems vs Optimization-Based Systems

While rule-based logic and Business Rule Management Systems (BRMS) offer straightforward solutions for simple if-then decision scenarios, they often fall short in the complex and dynamic environment of maintenance planning. This is because maintenance operations involve multiple, sometimes competing, objectives and constraints that change over time — conditions that rule-based systems are not designed to handle efficiently.

In contrast, optimization technology is specifically designed to navigate such complexity. It goes beyond static rules, dynamically adapting to new information and changing conditions to prescribe actions that best align with organizational goals. For instance, an optimization model can simultaneously consider factors like asset condition, availability of maintenance personnel, cost constraints, and the impact on operations to generate a maintenance schedule that minimizes downtime and costs.

Consider the example of a logistics company needing to optimize its vehicle maintenance schedule. A rule-based system might simply flag vehicles for service based on timespan and mileage. However, an optimization approach would consider additional factors like vehicle load, route terrain, and even predicted weather conditions to recommend maintenance only when really necessary, thus optimizing both vehicle uptime and maintenance resource utilization.

Furthermore, advancements in AI and machine learning technologies enrich these optimization models, enabling them to predict potential failures more accurately and adapt maintenance strategies in real-time based on evolving data patterns.

Two Examples: Subway Car Maintenance Planning & Technician Dispatching

Consider the complexities involved in maintaining subway train cars:

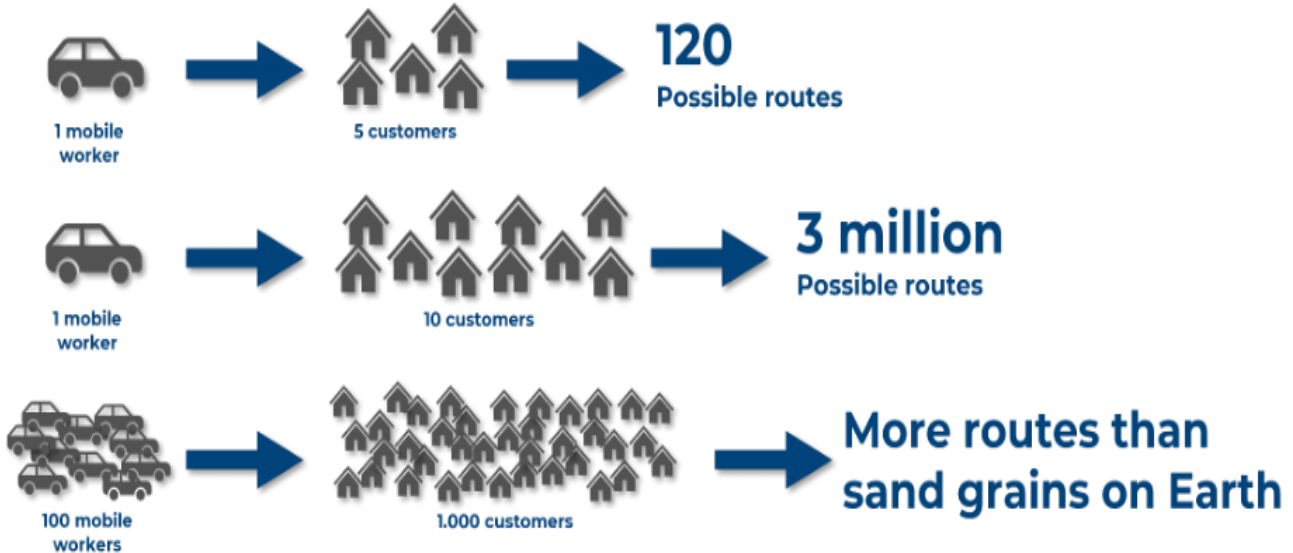
- **Asset Age and Condition:** Older cars require more frequent maintenance than newer ones.

- **Operating Environment:** Harsh environments (extreme heat, humidity, or salt) necessitate more frequent maintenance.
- **Spare Parts Availability:** Maintenance may need to be deferred if parts aren't readily available.
- **Balancing Cost and Efficiency:** If multiple service types are due soon, combining them can yield cost savings.
- **Staffing Availability:** Maintenance might need to be postponed or outsourced if there's a technician shortage.
- **Minimizing Service Disruption:** Schedule maintenance during times with the least passenger impact.

These various, interconnected factors create a complex scenario with countless plan possibilities. Traditional “if-then” logic-based software would not be able to generate optimal solutions in such situations.

Another classic use case for prescriptive analytics is optimizing technician dispatching and routing. Here are some factors to consider when scheduling mobile technicians:

- **Skillset and Certifications:** Matching technicians' qualifications to job requirements.
- **Technician Workload:** Ensuring schedules don't exceed available work hours.
- **Labor Compliance Regulations:** Adhering to work hour and day-off limitations.
- **Travel Efficiency:** Minimizing travel time and maximizing jobs completed per technician per day.
- **Job Prioritization:** Prioritizing urgent jobs based on their nature or client contracts.



Creating efficient schedules that address all these factors simultaneously becomes exponentially more difficult with increasing numbers of technicians or job sites.

These two examples highlight where **mathematical optimization**, a specific technique within prescriptive analytics, comes into play.

What is Optimization and How Does It Work?

Optimization is a branch of AI that comes out of operations research (OR). It is about selecting the optimal course of action from possibly millions or more possible options. The term “optimal” refers to the solution that solves for the **decision variables** in a way achieves the greatest satisfaction of a set of weighted **objectives** while taking into account various **constraints**.

For example, in the context of maintenance planning, the **decision variables** could be:

- The specific maintenance activities required for each asset.
- The scheduling of these activities, including target dates, durations, and frequencies.
- The assignment of tasks to specific technicians with the necessary skills and certifications.

The **objectives** guiding the optimization process are tailored to organizations' specific needs, such as:

- Minimize total maintenance cost over a defined period (e.g., 12 months).
- Minimize downtime for critical equipment that operates continuously.
- Maximize preventive maintenance coverage to reduce the risk of unexpected failures.

Constraints represent the limitations and real-world factors that the optimization process must consider:

- **Technician Travel Limitations:** Distance limitations on how far technicians can travel efficiently between jobs.
- **Maintenance Frequency Requirements:** Minimum or maximum time intervals between specific maintenance tasks for different assets.
- **Precedence Rules:** Defined order of operations that must be followed during maintenance shutdowns (e.g., replacing a specific component before another).

By incorporating these constraints, the optimization engine ensures the generated schedule is not only efficient but also feasible for your maintenance team to execute.

Optimization Applicability Across Maintenance Planning Functions

Optimization approaches are effective in addressing a wide variety of maintenance challenges. Here is a broader set of examples of where optimization can be applied across various maintenance functions:

- **Establishing Maintenance Policies:** Define preventive, predictive and corrective maintenance strategies aligned with your organizational goals.
- **Asset Uptime and Availability Targets:** Determine the optimal balance between uptime targets and associated maintenance costs.
- **Maintenance Frequency Optimization:** Identify the ideal frequency for each maintenance activity, maximizing equipment lifespan and minimizing unnecessary interventions.
- **Risk Management:** Evaluate the cost implications of accepting different levels of failure risk, facilitating informed risk-based decisions.
- **Staffing and Skills Management:** Optimize the number, skillsets and locations of maintenance workforce based on workload and asset distribution.
- **Spare Parts Inventory Management:** Determine optimal inventory levels and storage locations for spare parts, ensuring availability while minimizing carrying costs.

Maintenance Optimization Use Cases in Industry

Maintenance priorities can vary across different industries and types of industrial equipment. For example, some equipment may prioritize maintenance costs while others prioritize uptime or service level objectives. When a moving walkway at an airport requires maintenance, there may be some flexibility in waiting a day or more to ship parts and dispatch a repair technician, who can perform other types of maintenance while at the airport. However, when an aircraft part fails, it is urgent to restore it to service as quickly as possible. In such cases, all available resources are typically deployed with little regard for cost, since the lost revenue resulting from an aircraft not flying is substantial.

To better understand the varying maintenance objectives of different equipment assets, one approach is to segment them based on their usage patterns and type:




By using this segmentation, we can better understand the unique maintenance needs of each quadrant and develop tailored maintenance strategies to meet those needs. In the following sections, we'll explore specific use cases within each quadrant.

Heavy Equipment with Intermittent Usage - Aircraft Engine Lease and Maintenance Cost Optimization

A commercial airline, with a fleet of hundreds of single-aisle planes, faced the challenge of optimizing the 10-15-year long-term maintenance shop visit plan for owned and leased engines. The plan detailed retirement dates for owned aircraft as well as return dates for those that are leased.

The scope of the effort was to implement a decision support system that defines the optimal maintenance plan by minimizing overall maintenance and lease return compensation cash outflows while respecting operational, regulatory and contractual constraints.

Each engine is comprised of 20-30 life-limited parts (LLPs) with different life cycles, the task of defining a monthly plan that complies with all requirements and spans 120-180 months becomes exceedingly complex, with billions of possible configurations. Moreover, the plan requires frequent updates to accommodate unforeseen events.



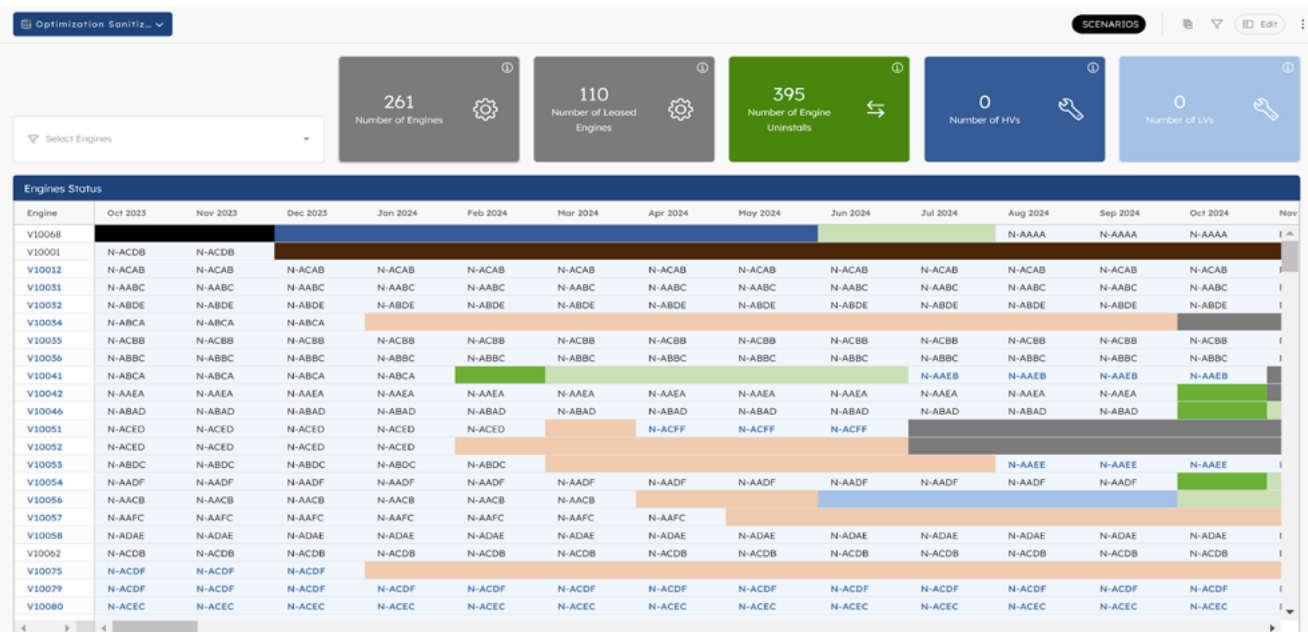
Given that maintenance shop visits can last several months and potentially exceed \$5 million in cost, a solution that streamlines the definition of the maintenance plan and optimizes cash outflows could yield substantial financial benefits.

DecisionBrain developed a custom-fit maintenance planning system, expected flight cycles and hours, maintenance history, leasing contract specificities and other operational and regulatory information, and generates maintenance plans at the engine level over a multi-year time horizon. The plans also indicate the scope of each maintenance visits, e.g., which limited life parts -LLPs- or groups of LLPs should be replaced.

The plan is created using powerful mathematical optimization solvers that produce fast results. By exploring billions of possible configurations, the software produces a plan that complies with operational and regulatory maintenance requirements, as well as contractual leasing characteristics while minimizing overall cash outflow.

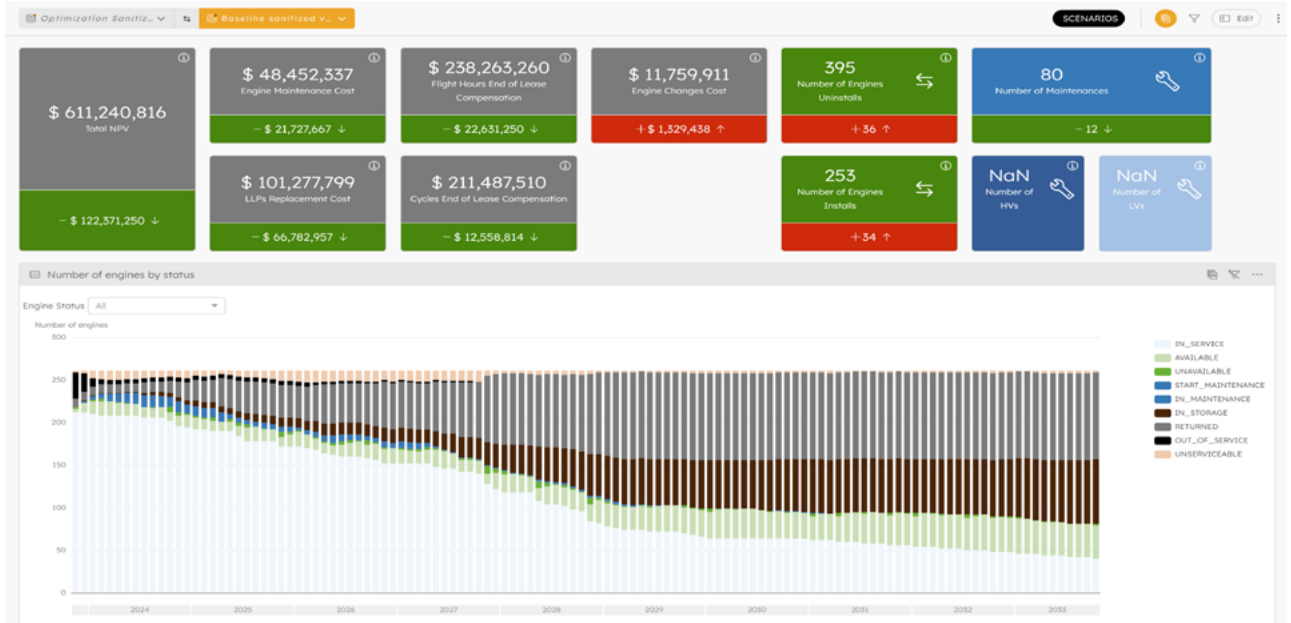
The system also allows planners to interact with the output. Additionally, to perform “what-if” analysis (e.g., simulate lease extensions, test different flight patterns, ...), assess the financial impact and define the most desired scenario.

Below is a screenshot of one of the application screens, which shows both KPIs on top as well as the auto-generated maintenance schedule represented in the gantt chart.



The Results

- Over 10% reduction in total maintenance and end-of-lease costs, leading to tens of USD millions in savings
- Possibility to quickly simulate the impact of changes in the fleet plan
- Enhanced compliance with the lessor and regulatory maintenance requirements
- Improved ability to plan maintenance during off-peak times, minimizing disruptions to flight schedules



Heavy Equipment with Continuous Usage - Aircraft Heavy Maintenance Optimization

In this use case, we focus on optimizing aircraft heavy maintenance, which is a complex and costly process that can take up to 60 days. Unlike the previous example where the planning horizon spanned a multi-year lease period, we make operational scheduling decisions on a shorter time horizon of 30-60 days. Heavy maintenance for aircraft is critical, and it's essential to complete all necessary activities during the maintenance period. However, the longer the aircraft is grounded, the greater the revenue loss. Thus, the objective is to minimize the maintenance duration while ensuring all necessary maintenance activities are completed.

This is a highly intricate scheduling problem that shares many similarities with manufacturing scheduling problems. It entails sequencing hundreds of maintenance activities while adhering to precedence constraints, which implies that not all tasks can be executed concurrently. In addition, just like manufacturing throughput is constrained by machine capacity and raw material availability, aircraft maintenance activities are also limited by material or part availability as well as workforce capacity, as they are performed by teams of technicians with varying skill sets.

To kick off the process, an initial maintenance plan is established using a range of data points, including the number of flight cycles and miles flown since the previous maintenance check, as well as any other sensor readings that might suggest the need for additional maintenance work.

As the maintenance process unfolds, the optimization system constantly evaluates the latest data and information, and then generates updated maintenance plans on a daily basis. This way, any new findings or insights that are collected by technicians during the day can be quickly factored into the optimization process, ensuring that the maintenance plan remains as up-to-date and accurate as possible.

The proof-of-concept for this project demonstrated potential savings of several million dollars per year, offering a rapid and substantial return on investment.

This maintenance optimization approach is not limited to aviation but can be applied across various industries in which maintenance must be performed on heavy equipment that is continuously used, such as pulp and paper mills, providing companies with the means to improve asset utilization and minimize downtime costs.

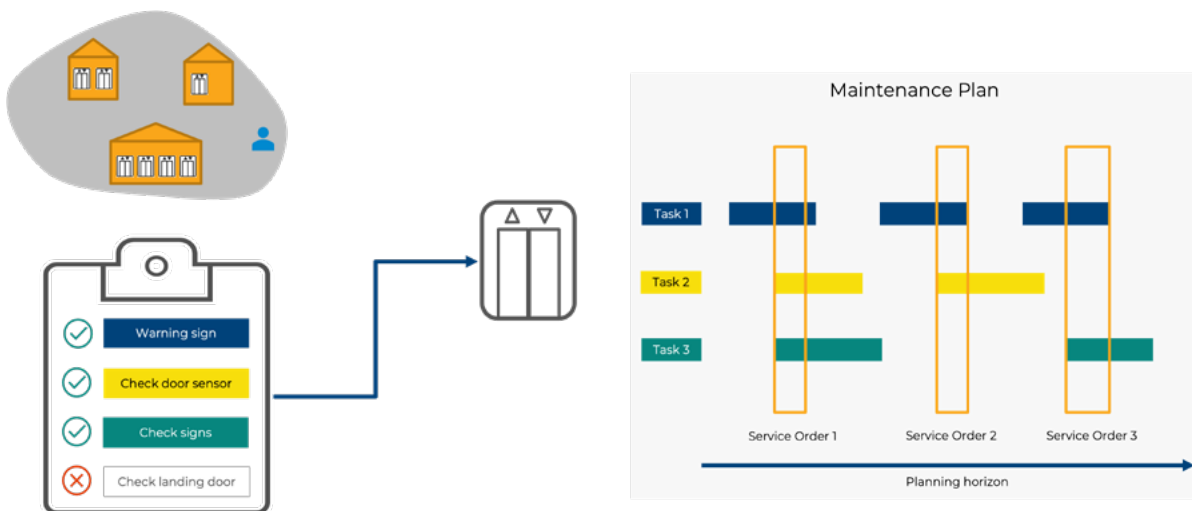
Light Equipment with Intermittent Usage – Dynamic Tactical Elevator Maintenance Planning

A manufacturer of elevators and escalators sought to improve their maintenance planning process by moving away from frequency-based scheduling and adopting a more data-driven, outcome-based approach. They recognized that not all units in all locations experience the same rate of performance degradation or are used at the same frequency, and they wanted to optimize their maintenance plans accordingly.

To achieve this, the company identified several factors that could influence maintenance planning, including usage, age/type/model, and environmental factors. Specifically:

- **Usage:** The number of times the elevator is used per day, week, or month can impact how often maintenance is required. Door open/close cycles and floors served can also be a factor.
- **Age/Type/Model:** Older elevators may require more frequent maintenance than newer ones or they may require different skillsets or tools. Hydraulic elevators require more maintenance than traction models.
- **Environmental factors:** For example, elevators located in areas with high humidity or extreme temperatures may require more frequent maintenance.

The company wanted to create optimized tactical plans over a rolling 12-month time horizon, leveraging the most current asset usage information to create and regularly revise their plans. Since the elevator units were geographically distributed, location and travel time for technicians was a consideration and minimizing travel time was a key performance indicator. One efficiency improvement opportunity they identified was to create smart service “bundles,” grouping a set of maintenance activities that were closely located and all coming due for maintenance around the same time into a smart bundle that could be carried out by technicians with the right qualifications. (See the below figure). The plans also had to consider workforce capacity by week and by skillset.



The company partnered with DecisionBrain to implement an optimization-powered dynamic maintenance planning system and initially rolled it out to a test market in Europe. The results were positive, demonstrating significant improvements in maintenance efficiency and planning effort.

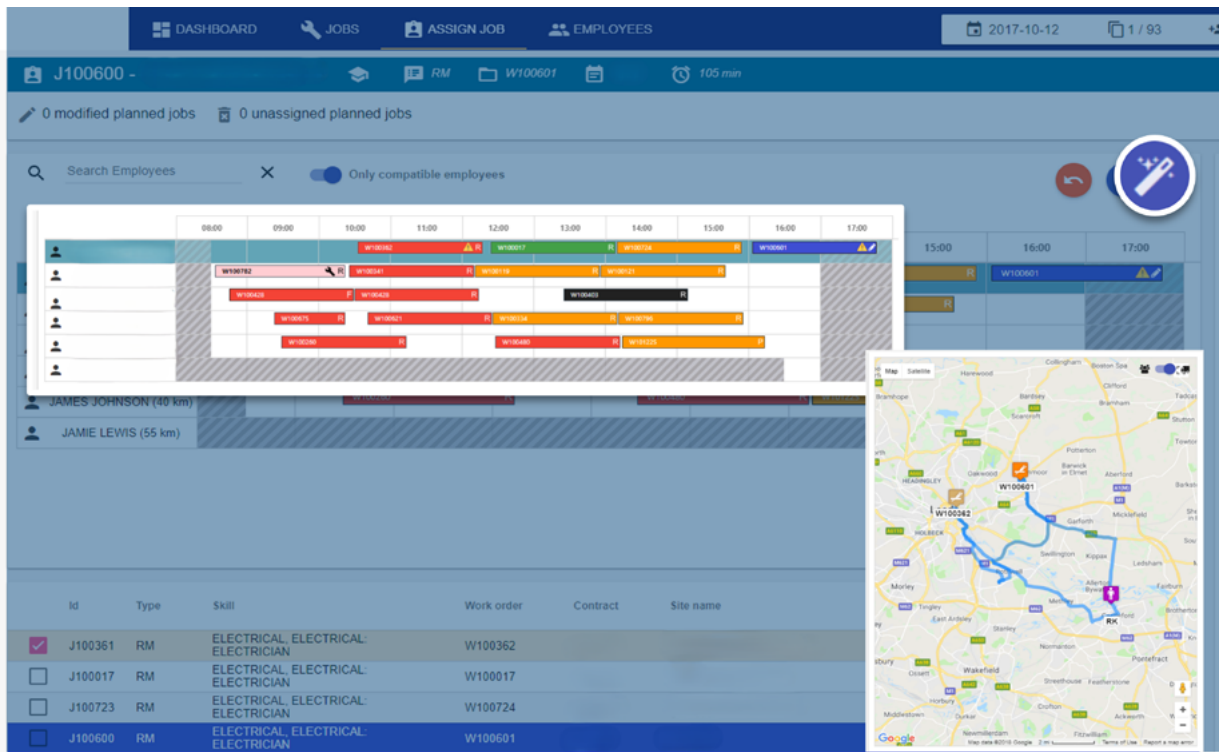
Before rolling the solution out more broadly, the company identified several additional considerations, including renegotiating contracts with clients to agree to outcome-based (or uptime-based) service contracts rather than frequency-based contracts.

Future projects are focused on expanding the use of the optimization-powered system to operational scheduling. One key consideration is incorporating employee preferences into the planning process to improve job satisfaction. To this end, the software enables technicians to signal their availability for the next job via their mobile devices, and select their preferred criteria for job allocation, such as job priority, efficiency (measured in terms of the number of jobs completed), or “Eco” (minimizing travel distance). The system then provides technicians with a list of compatible jobs based on their qualifications, allowing them to choose which job to accept. By empowering technicians to select their preferred jobs, this approach can increase employee satisfaction and engagement, leading to further improvements in maintenance efficiency.

Light Equipment with Intermittent Usage – Facility Maintenance Technician Dispatch and Scheduling

A UK-based Mechanical and Engineering (M&E) service provider faced a challenge with their weekly assignment and scheduling of jobs for their technicians, who were responsible for providing both Planned Preventive and Reactive Maintenance services for over 1,600 commercial clients across 60,000 locations in the UK. The company wanted to increase productivity of their technicians in the field and reduce travel and idle time.

The company partnered with DecisionBrain to implement a Field Service Optimization (FSO) solution. The FSO system was customized to meet the needs of the various organizational stakeholders, from the project executives, planners, group managers and field service workers. FSO, provided several benefits, including the ability to quickly focus on urgent issues, adapt and adjust plans to unexpected events, and more easily collaborate across teams.



After implementing the scheduling optimization system, the company experienced a significant increase in technician productivity, with technicians completing an average of four jobs per technician per day, compared to 2.4 jobs per technician per day prior to the optimization. This improvement was attributed to several factors, including:

- Better routing of technicians, reducing travel time between jobs.
- Improved matching of jobs with technicians possessing the right skills.
- Reactivity to unexpected events, such as jobs taking longer than expected, by performing a local re-optimization of the schedule.
- Opportunistic adjustment of planned activity schedules to take advantage of situations when a qualified technician was called to a nearby job earlier than the scheduled maintenance.

By incorporating an optimization-powered scheduling system, the company was able to maximize the productivity of their technicians in the field while reducing travel and idle time. This approach can be applied to other service industries with field technicians, enabling companies to improve their scheduling and dispatching processes, and increase efficiency and profitability.

Conclusion

The availability of data, whether from IoT, telematics and sensors or other situational or forecast data about our business, has opened up tremendous new possibilities to optimize the way we plan and execute maintenance activities.

Through analyzing usage patterns, sensor output, and other potentially correlated or causal factors of failures, we can build a model to calculate the probability of failure over time. By joining these predictive models with details about failure impact and maintenance effort/cost, we can define the right maintenance policies. Leveraging AI and optimization-powered software, we can then build optimized plans and schedules based on the latest data, goals, and constraints.

With dynamic optimization-powered scheduling software, we can quickly react to surprises such as breakdowns, material backlogs, or workforce shortages. Additionally, by simulating different policies or situational scenarios, we can understand the impact of different decisions on key performance indicators, such as service cost, efficiency, service level, workload balance, and more.

By leveraging AI techniques to carry out these steps, including machine learning and mathematical optimization, we can prepare, plan, schedule, simulate, and ultimately succeed in achieving our desired financial and service level outcomes.

[Contact us](#) if you have any questions or want to discuss further!

