Predictive Maintenance and an integrated data platform as drivers of efficiency in the pharmaceutical industry

Last updated: 08/2026

zeiss.com/digital-innovation

ZEISS

Seeing **beyond**

This image was created with the help of Generative AI

Introduction

In the pharmaceutical manufacturing environment—where GMP-compliant production processes, highly specialized equipment, and strictly regulated process windows converge—technical malfunctions can have far-reaching consequences. Deviations in critical plant components not only lead to production interruptions but also potentially to out-of-specification results, time-consuming deviation investigations, repeat qualification or validation measures, delayed batch release, and, in the worst case, the loss of entire production batches. As a result, the maintenance of machines and production equipment becomes a key driver of product quality, batch yield, delivery capability, and total cost of ownership. At the same time, conservative, rigid maintenance intervals often lead to over-maintenance, the unnecessary replacement of expensive wear parts, and inefficient use of personnel resources.

Against this backdrop, predictive maintenance is gaining increasing strategic importance in the pharmaceutical industry. It transforms traditional preventive maintenance measures for production facilities into a process of data-driven forecasting of the remaining service life of relevant components and equipment. Based on extensive process and machine data—such as sensor readings from granulators and fluidized-bed systems, pressure and temperature parameters from autoclaves, data from cleanroom monitoring systems, deviation and event logs, as well as MES, LIMS, and ERP information—maintenance measures can be planned on a condition-based basis, calibration and maintenance intervals can be dynamically adjusted, and unplanned equipment failures can be significantly reduced.

Compound annual growth rate of the global predictive maintenance market through 2034

24.3%[1]

This white paper examines how manufacturing companies in the pharmaceutical industry can reduce their direct and indirect maintenance costs, minimize unplanned production downtime, and improve the Overall Equipment Effectiveness (OEE) of their equipment through the use of predictive maintenance. The focus is on both process and economic metrics as well as the role of a central data platform as the technical foundation for implementing predictive maintenance scenarios in highly regulated pharmaceutical production environments, taking data integrity and validation requirements into account.

1 Source: FortuneBusinessInsights (2026): Predictive Maintenance Market Size, in: https://www.fortunebusinessinsights.com/predictive-maintenance-market-102104 (23.01.2026)

Table of contents

1. What is predictive maintenance?

Maintenance is necessary to operate machinery and equipment in a safe condition. According to DIN 31051 (Fundamentals of Maintenance), maintenance is divided into four categories: inspection, preventive maintenance, corrective maintenance, and improvement. Maintenance is defined as a measure to delay the depletion of existing wear reserves and is distinguished from repair, which is defined as a measure carried out to restore the function of a defective object. Figure 1 illustrates these relationships.

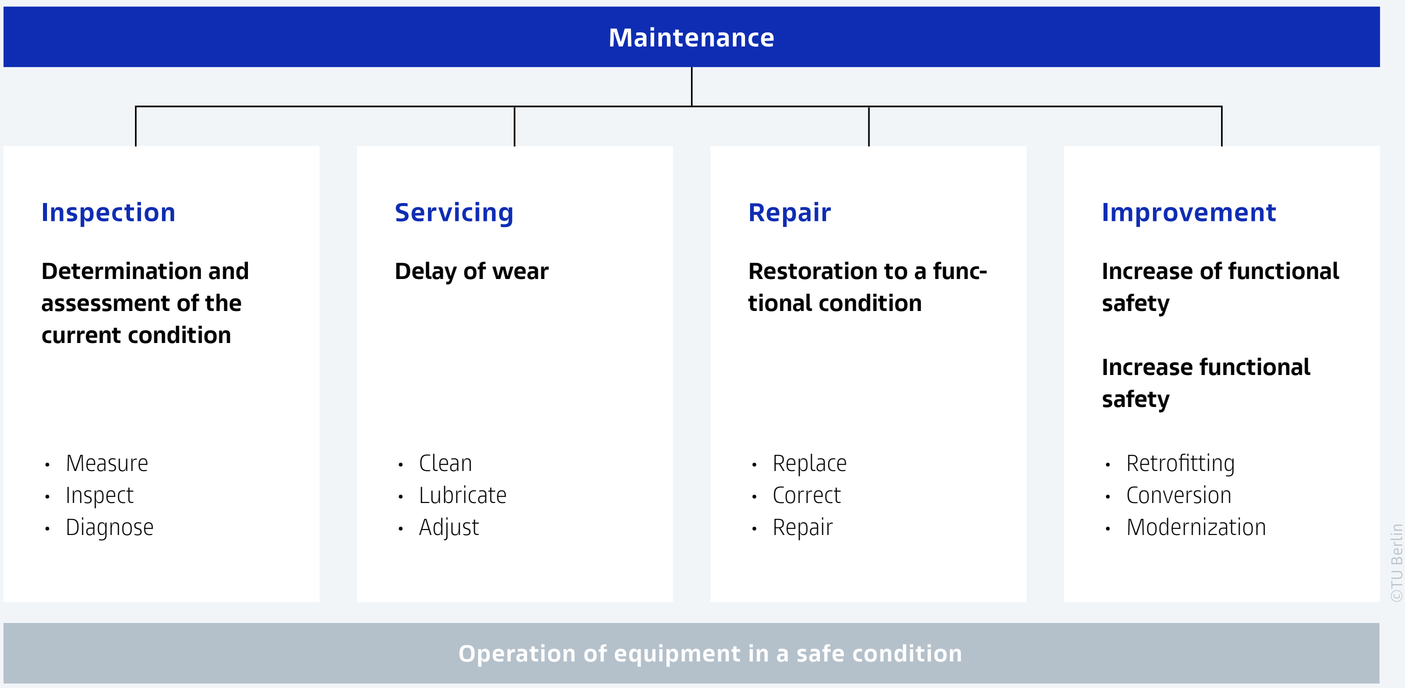

Figure 1: Maintenance measures according to DIN 31051

However, DIN EN 13306 (Maintenance—Maintenance Terminology) does not recognize the term "maintenance" and defines only three different types of maintenance: preventive maintenance, corrective maintenance, and improvement, as shown in Figure 2. In this context, the measure "maintenance" correlates most closely with the type of preventive maintenance. The technical maintenance measure "inspection" is no longer listed separately here but is a natural part of the three types of maintenance. While inspection describes the process by which a person verifies the conformity of the relevant characteristics of an object (system) (using measurement, observation, or testing), condition monitoring is distinguished by the fact that it serves to determine any changes in the object's parameters over time. Monitoring can take place continuously, at regular intervals, or after a specified number of operating cycles, and is typically performed while the system is in operation.

Maintenance based on condition monitoring is called condition-based maintenance. If it is possible to use algorithms based on condition monitoring to predict the remaining service life, this approach is referred to as predictive maintenance (according to DIN EN 13306) or, synonymously, proactive maintenance.

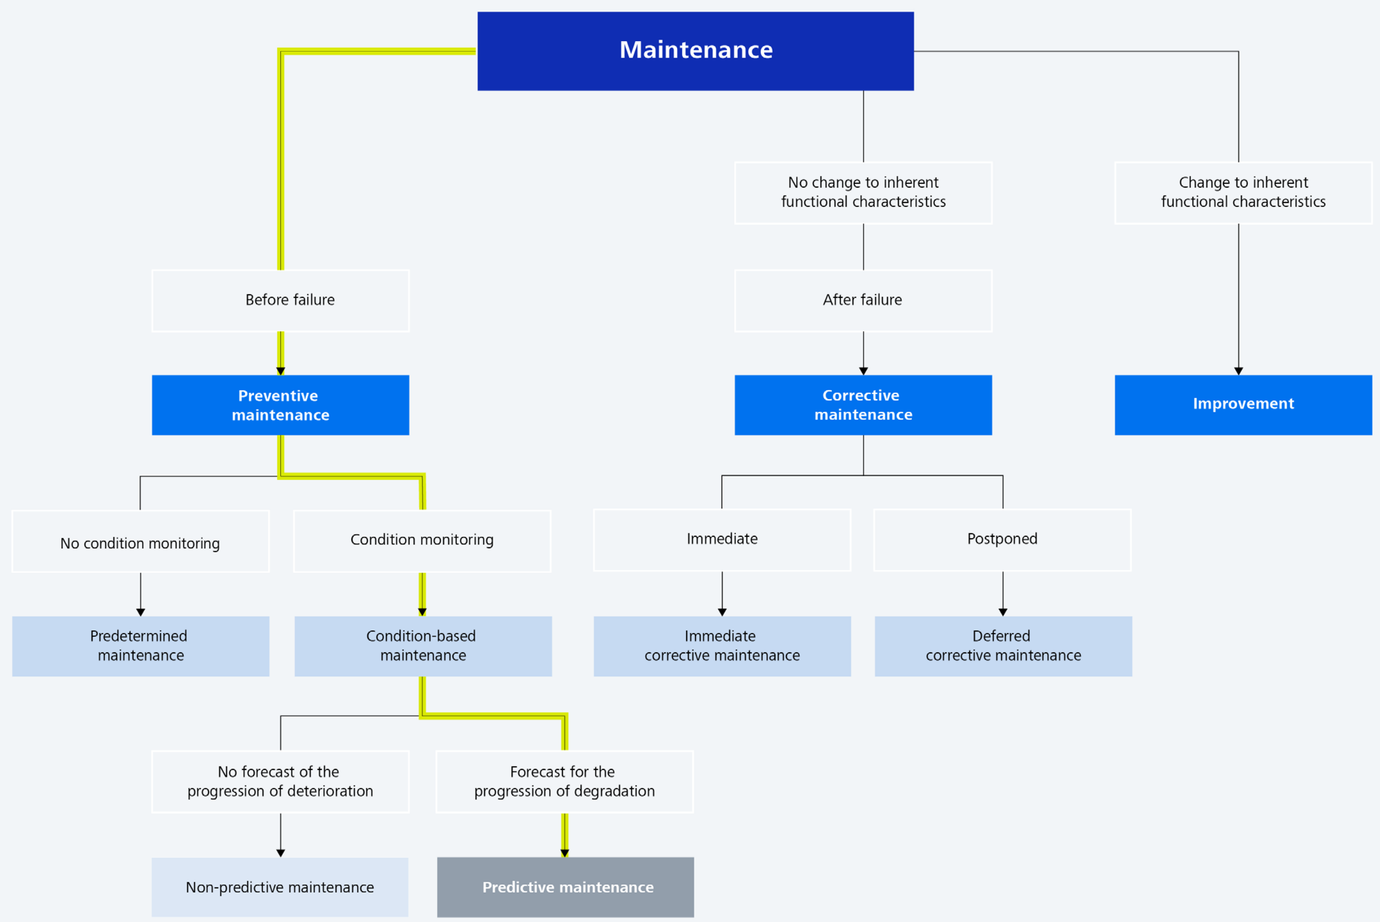

Figure 2: Types of maintenance according to DIN EN 13306

When the relationships are simplified for practical application, three focus areas emerge that together constitute predictive maintenance: organizational and planning measures, condition monitoring of the equipment, and algorithm-based forecasting. These three focus areas are by no means independent of one another. Instead, they are closely interconnected and interact continuously. For example, condition monitoring enables identification of which parts of the equipment will require maintenance during the next maintenance cycle. This insight directly triggers concrete planning actions, such as spare part provisioning.

Approaches for establishing a unified data foundation are discussed in the white paper “Industrial Data Platform”.

Many companies already use software support for partial solutions. Maintenance work orders are often managed through work order management systems that may be integrated into ERP platforms. In addition, standalone solutions for condition monitoring or failure prediction of individual assets are commonly used. The challenge in evolving toward end-to-end predictive maintenance lies in the fact that these isolated solutions do not interlock. As a result, value creation remains limited to the original scope of each individual tool.
To unlock the full potential of predictive maintenance, solutions covering all three focus areas should ideally be implemented on a shared software infrastructure. This approach avoids data discontinuities and enables automated workflows across the entire maintenance process chain.

Approaches for establishing a unified data foundation are discussed in the white paper “Industrial Data Platform”.

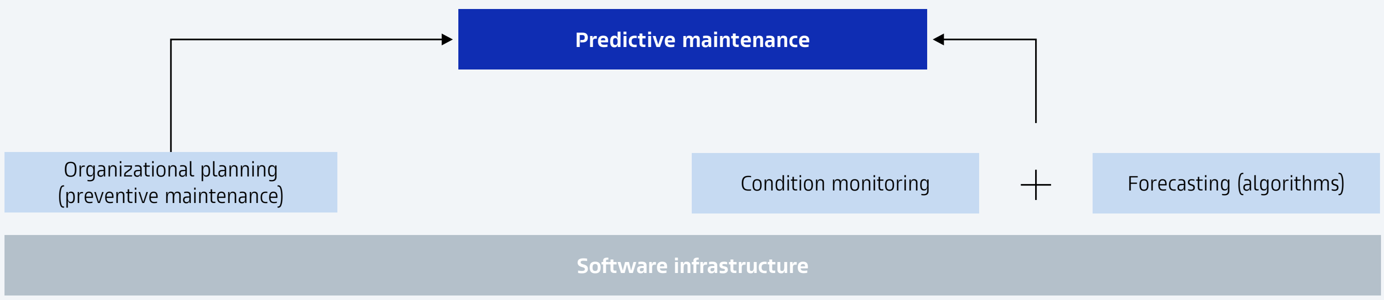

Figure 3: Key focus areas of predictive maintenance

2. Process objectives of predictive maintenance

The introduction of predictive maintenance primarily aims to reduce direct maintenance-related costs. These include labor costs for maintenance activities, material costs for spare parts and wear parts, and expenses for external service providers. In addition, predictive maintenance can significantly reduce indirect costs, such as production losses caused by mainte- nance-
related downtime, quality losses resulting from delayed corrective actions, and capacity losses during unexpected equipment failures.

Two fundamental process objectives drive the implementation of predictive maintenance. The first objective is to reduce the share of unplanned maintenance activities. The second objective is to extend maintenance cycles that are currently based on predetermined intervals. Both objectives contribute directly to lowering total maintenance-related costs.

Reduction of the proportion of unplanned maintenance

In practice, equipment failures typically require immediate corrective maintenance to resume production. In such situations, the available lead time for planning and preparation is effectively zero. As a result, technical, logistical, and administrative delays directly increase equipment downtime.
Preventive maintenance, by contrast, enables evaluation and preparation while the equipment is still operational. Potential delays can be identified and mitigated through structured planning prior to execution. This significantly reduces non-functional time and improves overall maintenance efficiency. Figure 4 compares preventive and corrective maintenance.

If no replacement is available for a failed tool, or if the tool is tightly coupled within a production line, an unplanned shutdown may bring the entire line to a standstill. The resulting production losses disproportionately increase unit manufacturing costs across the factory.

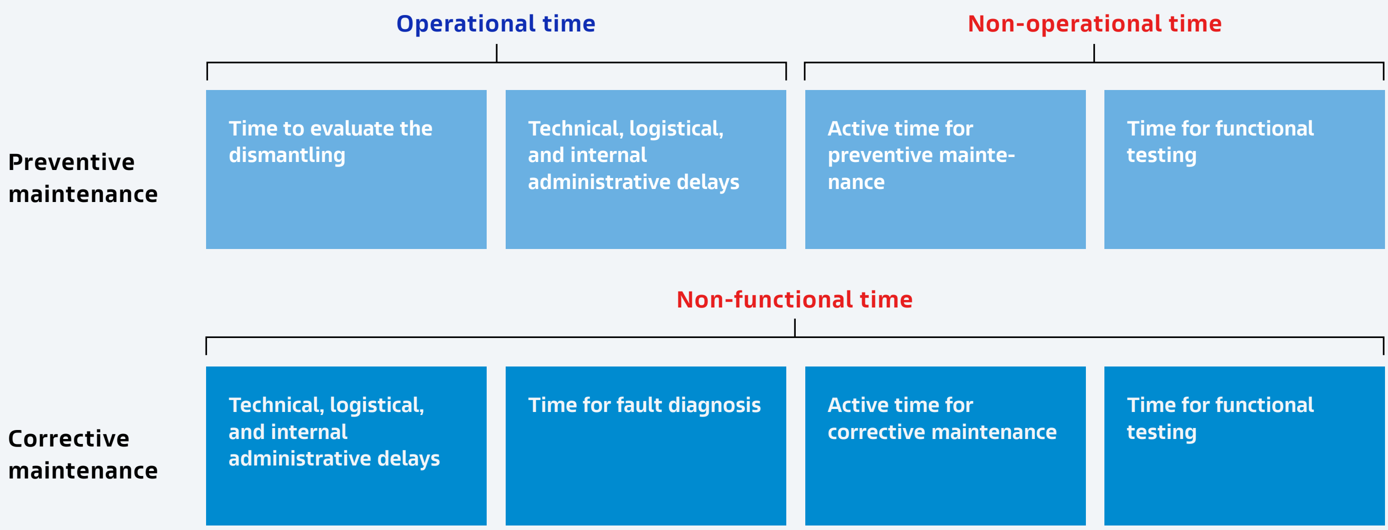

Figure 4: Comparison of preventive and corrective maintenance

Extension of predetermined maintenance cycles

To avoid unplanned downtime, many companies rely on fixed maintenance intervals that are based on manufacturer recommendations or internal operational experience. Figure 5 compares a preventive maintenance cycle with the theoretical maximum time span until maintenance is required over the progression of asset wear.

Because failure risk typically increases in a nonlinear manner over time, these intervals are often defined more conservatively than technically necessary. As a result, maintenance activities are performed more frequently than required, which leads to inefficient resource utilization, premature replacement of components, and elevated labor costs.

Simply extending maintenance intervals based on empirical assumptions is often not feasible, because the maximum achievable interval depends on specific fault modes and operating conditions. In many organizations, maintenance intervals have already been optimized to the greatest extent possible within acceptable risk limits.

The appropriate response is a strategic shift toward condition-based and forecast-driven maintenance. Under this approach, maintenance cycles are no longer static. Instead, they are continuously adjusted based on the actual condition of the equipment and its predicted degradation behavior. This minimizes unnecessary maintenance effort while avoiding an increase in operational risk. Figure 6 illustrates the difference between predetermined maintenance cycles and predictive maintenance cycles.

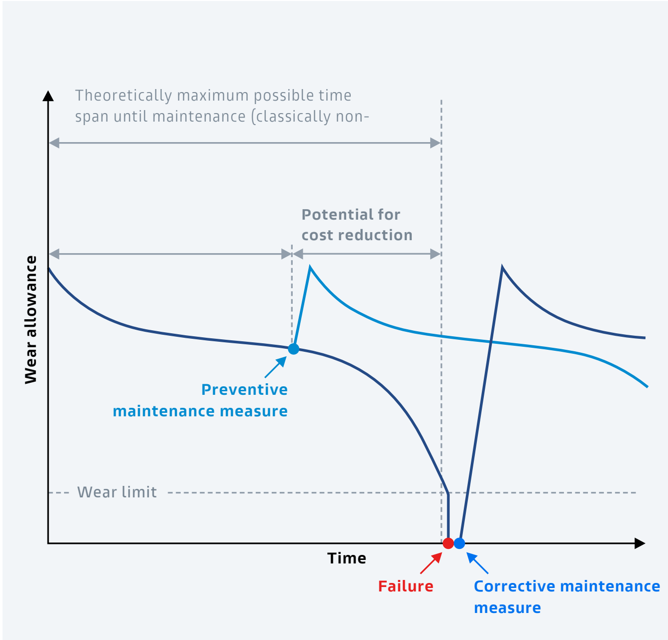

Figure 5: Diagram showing the progression of the wear allowance over time

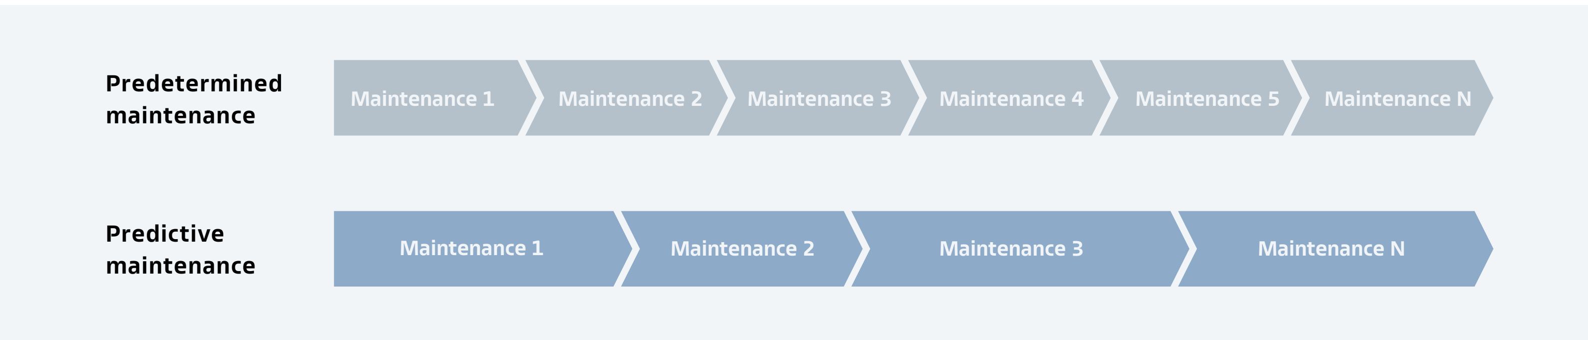

Figure 6: Comparison of maintenance cycles in predetermined and predictive maintenance

3. Digitalization strategy for a typical maintenance scenario

Maintenance workflows are well understood across industries and differ primarily in their implementation rather than in their fundamental approach. To ensure a smooth and reliable process, several coordinated measures and process steps are required.

The workflow begins with the definition of a work order. This is followed by work order documentation that captures all relevant information in a structured form. An analysis of the documented content helps clarify the specific maintenance requirements. Based on this analysis, a binding maintenance plan is created and tailored to the needs of the organization. The plan specifies the location, scheduled date, required maintenance activities, and the characteristic values to be monitored. Preparation for execution includes providing the necessary resources as well as implementing preliminary measures such as appropriate workplace equipment and safety devices. Before maintenance activities begin, the preparation is reviewed, and execution is formally approved. After the maintenance work has been completed, a functional test is performed to verify proper operation. Finally, feedback is recorded to document the process and identify opportunities for improvement in future maintenance activities.

If a company aims to digitalize the full process, it should not start with the most complex use case, forecasting. Instead, predictive maintenance should be implemented step by step by breaking down the overall objective into manageable digitalization packages.

Figure 7: General phases of repair

3.1 RAG-Based knowledge assistance for maintenance personnel

Technical analysis requires a solid understanding of equipment functionality and manufacturer-recommended procedures, which are typically documented in manuals and operating instructions. Due to the large volume of documentation, locating and extracting relevant information from these sources is often time-consuming for technicians.

Retrieval augmented generation (RAG) combines language models with document-based knowledge sources. Technicians can submit natural language requests, such as generating an illustrated step-by-step repair guide for a specific component. The system automatically compiles the relevant information and delivers it in a structured format. This significantly accelerates knowledge acquisition and supports consistent execution across different sites.
Table 1 provides an overview of this use case.

Overview of maintenance technician guidance using a RAG system	
Phase of the maintenance process	• Analysis of the order content and planning of the characteristic values to be observed
Required data	• Order description • Operating instructions, repair manuals, etc. in the form of electronic documents (e.g. PDF)
Required software infrastructure	• Can be created independently of other use cases in a wide variety of software technologies
Approximate time frame for implementation	• Days to a few weeks
Key figures and added value for determining the return on investment (ROI)	• Maintenance time per case is reduced • The resulting documentation is available in multiple languages

3.2 Digital fault and action management

Until predictive maintenance is implemented across all maintenance cases, unexpected faults will continue to occur and will require corrective maintenance. A digital fault and action management system supports many of the required processes to minimize downtime. The system fulfills three core functions.

Fault capture
Fault and alarm messages from machines are typically captured automatically. In addition, employees can describe the fault in greater detail using guided software dialogs, for example, on a tablet. Because shop floor employees are not technical authors, simplicity and clarity of the input process are critical. A direct benefit of systematic fault capture is the ability to generate meaningful production metrics, such as overall equipment effectiveness (OEE) and Pareto analyses.

Action recommendation system
Based on captured fault data, the action recommendation system identifies the corrective action that has most efficiently resolved comparable issues in the past. By transferring digitized operational experience to the responsible team, it reduces the impact of staff turnover and supports effective knowledge transfer across shifts and sites. To ensure data quality and avoid low-value content, the recommendation system requires editorial governance.

Action reporting system
To automatically determine which measures are effective in real-world conditions, employees must document the actions they have taken and provide feedback on whether the proposed measures have achieved the desired outcomes. At the same time, data from other sources, such as time recording systems, can be used to evaluate the measure quantitatively.
The reporting component feeds results back into the recommendation logic to keep it up to date.

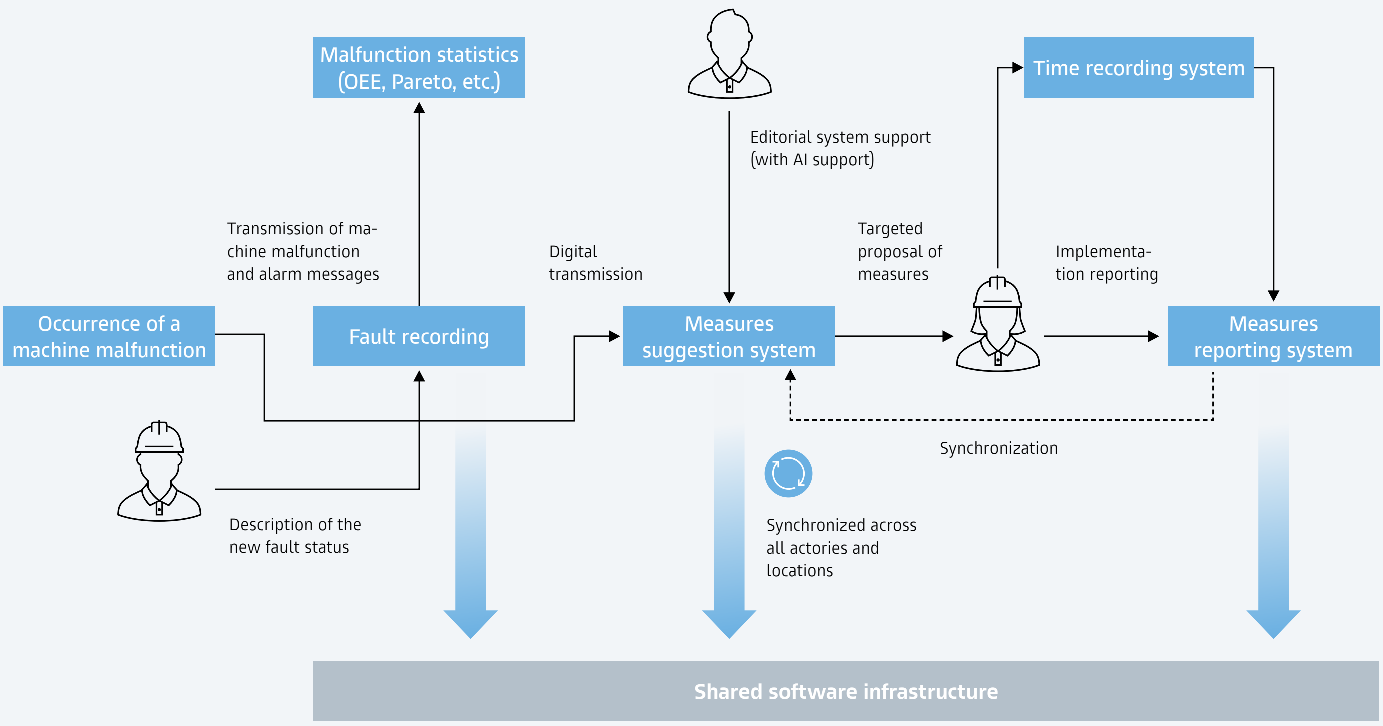

There are virtually no limits to integrating a fault and action management system with other business processes. For example, based on the recommended action, required spare or wear parts can be automatically ordered or reserved. The system can verify whether the person responsible has the necessary qualifications and immediately trigger a maintenance work order. In addition, a forecast of the expected production line downtime can be generated in real time, providing a valuable basis for production management decisions. At higher maturity levels, fault and action management systems also deliver high-quality data that can be leveraged for advanced forecasting models.

Figure 9 illustrates an example of software for fault and action management.

Overview of digital fault and action management	
Phase of the maintenance process	• Analysis of the order content and planning of the characteristic values to be observed • Feedback and documentation
Required data	• Fault messages from machines • User input by employees
Required software infrastructure	• All subsystems should use the same infrastructure. For cross-location systems, a cloud infrastructure is advantageous.
Benchmark for implementation	• Weeks to a few months
Key figures and added value for determining ROI	• Digitalized knowledge management enables employees to familiarize themselves with the problem at hand much more quickly. > ROI can only be determined meaningfully on the basis of overall production performance.

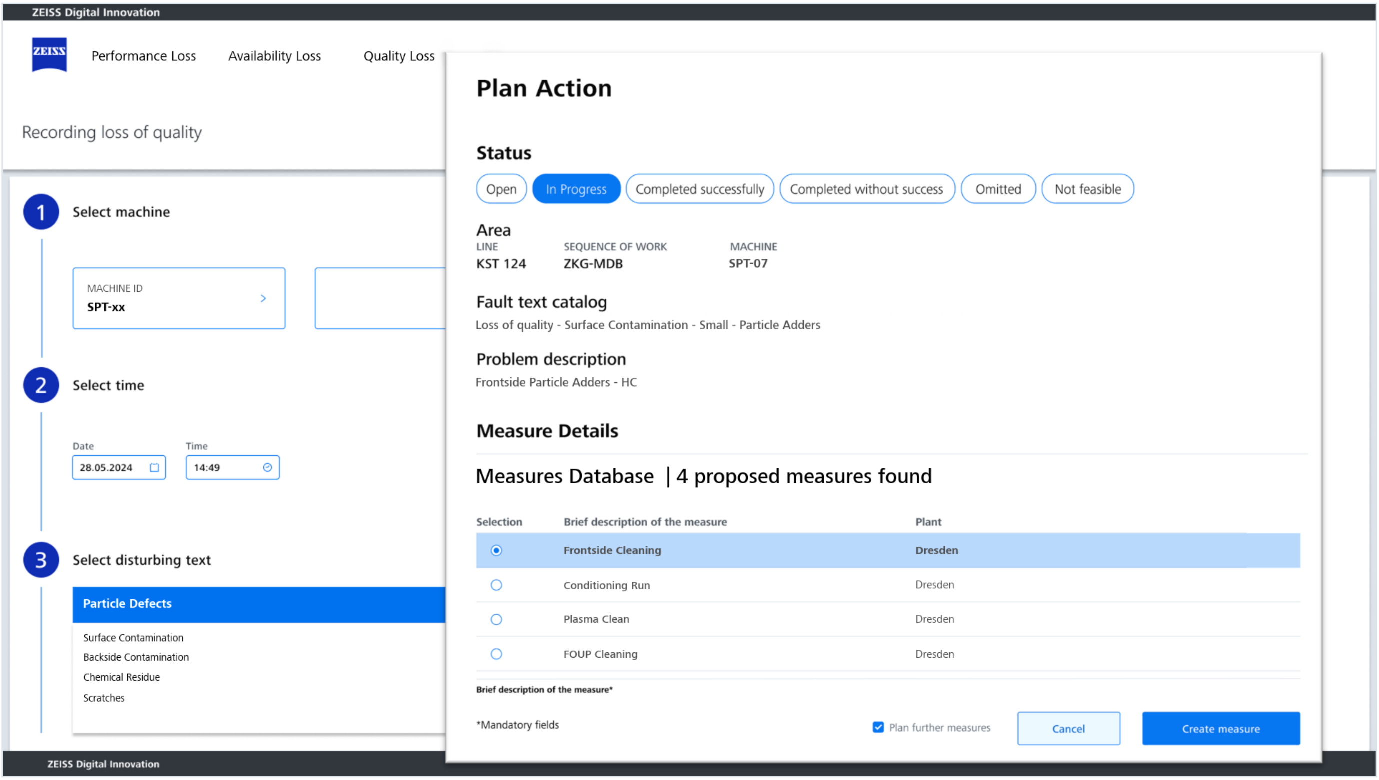

Figure 9: Example of software for fault and action management

3.3. Implementation of the predictive maintenance process

Within the digital fault and action management use case, a digital data source has already been established in the form of fault and warning messages generated by machines. Figure 10 provides a structured overview of the data sources that can and should be leveraged for predictive maintenance.
Data sources already available in production environments include the machine itself with its integrated sensors, as well as company-wide planning and controlling systems. Because many machines were not originally designed for predictive maintenance use cases, additional sensors may need to be retrofitted. These sensors are typically not integrated into the machine control system and therefore must be evaluated separately. Common examples include vibration sensors and camera systems combined with image analysis software for optical condition monitoring. Another important category is environmental sensing. Although it could be considered a form of external sensing, it is treated separately because it does not map one-to-one to a specific machine. Typical examples include temperature, humidity, and air pressure sensors, which are often available through cleanroom or facility management systems.

Even when much of this data already exists within the organization, it is often stored in isolated silos and cannot be accessed in a unified or systematic way. The first step toward predictive maintenance is therefore data enablement, with the target state being shared, structured, and consistent access to all required data. The original conceptual illustration introducing predictive maintenance (Figure 3) also implies an additional data source that has not yet been explicitly addressed: preventive maintenance planning. This planning should be supported by software solutions that provide structured and accessible data. Basic spreadsheet-based planning is less suitable due to a higher risk of errors, limited collaboration capabilities, and the lack of robust version control, all of which negatively affect planning quality. To achieve reliable and scalable planning results, a standard or custom database-driven software solution is therefore required.

Company-wide planning and controlling systems
- ERP
- MES
- Shift planning software
- Warehouse management software
- Energy measurement systems
- Documentation systems
- Employee qualification matrix
- Personnel management software
- ...

Machine-integrated sensor technology
- Part of the machine supplied by the manufacturer
- Typically connected to the machine control system
- Status data is reported to external systems via a communication channel of the machine control system (e.g., OPC UA)

Environmental sensors
Additional sensors installed in the production hall to measure environmental data such as temperature, humidity, and illuminance

Additional external sensors
Additional sensors installed by the operator that measure more data than is possible with the machine-integrated sensors (e.g., vibration). Typically not connected to the machine control system and must be read and processed separately. Camera systems are part of this category.

Figure 10: Classification of data in the context of predictive maintenance

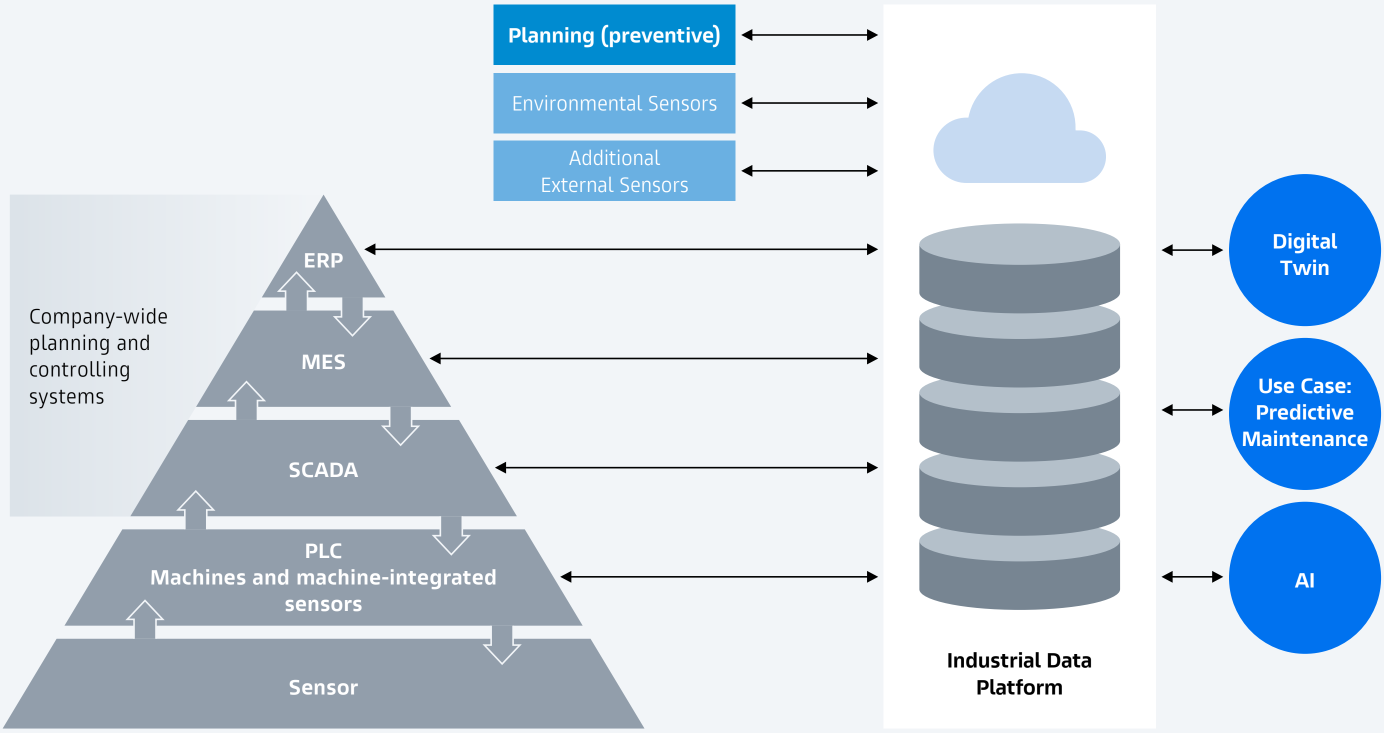

Figure 11: Data enablement using predictive maintenance as an example

When the logical view of predictive maintenance is translated into a data-flow perspective, the distinction between data enablement and value creation through forecasting and visualization becomes clear, as illustrated in Figure 11. Forecasting algorithms access a shared data repository, while visualization components consume and present the evaluations and results generated by the forecasting layer.

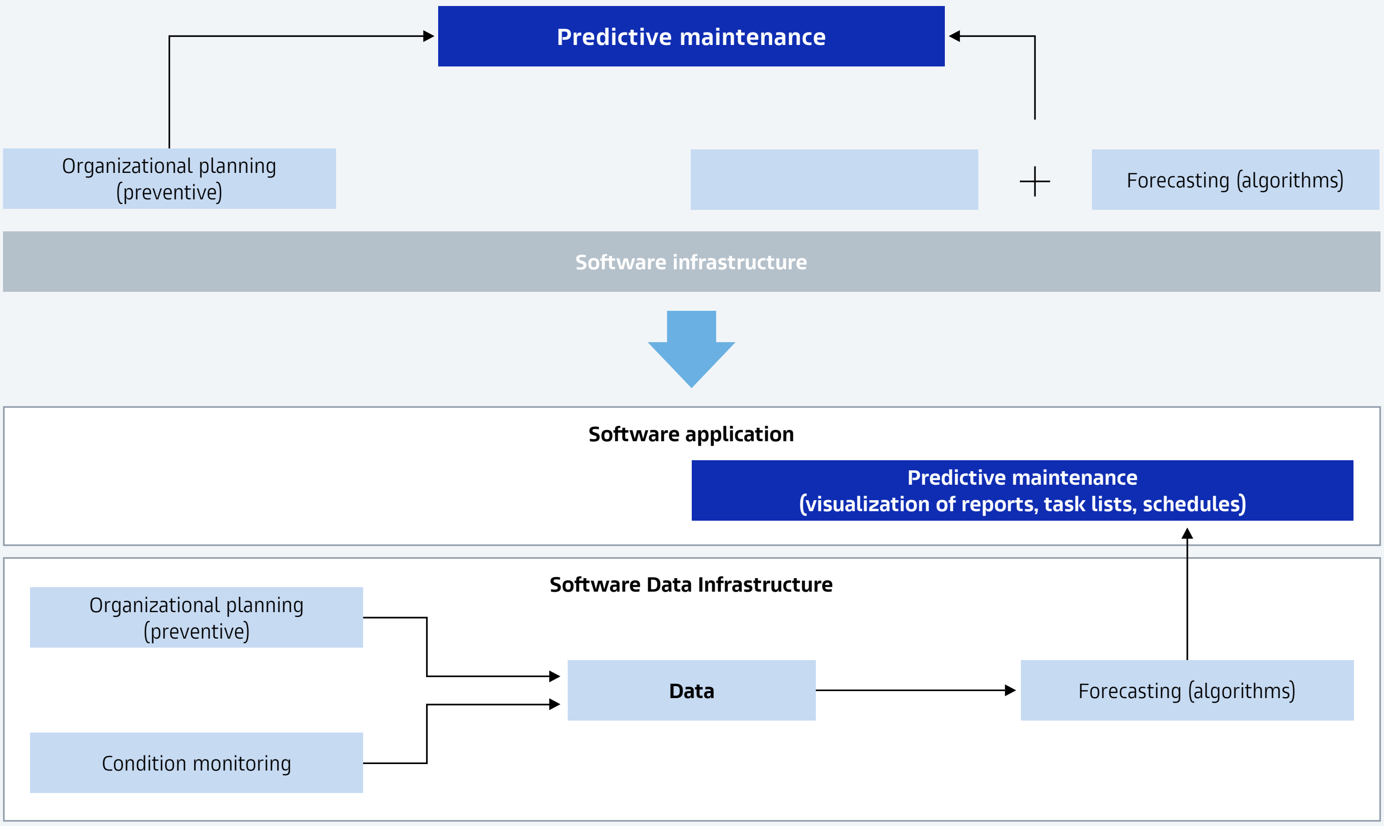

Figure 12: Foundations of predictive maintenance: transformation into a data-flow view

In general, forecasting algorithms can be understood as an optimization problem that combines remaining useful life estimation with preventive maintenance planning. Among these two aspects, remaining useful life estimation is the more demanding task. A combination of descriptive statistics and exploratory data analysis provides a solid foundation for subsequent machine learning approaches. It is advisable to establish data models at both the component level and the equipment level and to extend these models with dedicated analysis models. Component manufacturers, such as pump suppliers, may already provide analysis models. However, these models are often accessible only within the manufacturer's own infrastructure and software ecosystem. This creates challenges for operators, including increased contracting complexity across multiple suppliers and, in the worst case, the need to manage numerous software interfaces. A strategically stronger approach is therefore to operate both data models and analysis models on the operator's own infrastructure.

Analysis models for complete machines are currently limited in availability and, therefore, often must be developed by operators themselves. To enable rapid scaling, standardization of methods is essential so that value created once can be efficiently transferred to additional machines and use cases.

The visualization of algorithmic results can be tailored to the specific needs of an organization and standardized across the enterprise. Suitable formats include reports and dashboards, task and instruction lists for employees, or fully automated production schedules and maintenance plans. These views can be adapted to the production environment and configured according to user roles, balancing flexible, role-specific perspectives with standardized layouts and a consistent look and feel.

Predictive maintenance	
Phase of the maintenance process	• Analysis of the order content and planning of the characteristic values to be observed • Maintenance schedule (location, date, measures and characteristic values to be observed) • Preparation for implementation • Review of preparation and preliminary measures and approval • Feedback and documentation
Required data	• Company-wide planning and controlling systems • Machine-integrated sensor technology • Additional external sensor technology • Environmental sensors • Planning data
Required software infrastructure	• The software and data infrastructure should be based on the principle of the industrial data platform
Target value for implementation	• Months
Key figures and added value for determining ROI	• Predictive maintenance extends maintenance cycles and thus prevents waste of resources. ROI through reduced maintenance costs.

4. Return on investment for predictive maintenance

An ROI calculation can be performed for each individual measure based on defined key performance indicators (KPIs). DIN EN 15341, titled "Maintenance, Key Performance Indicators for Maintenance", defines a total of 183 KPIs, which can be weighted according to their economic relevance for a given organization. Measuring and evaluating these KPIs is, however, highly complex and typically requires dedicated software support. As a result, many companies are unable to reliably quantify their current performance. Consequently, potential future improvements cannot be evaluated on a sound cost basis. This lack of transparency makes it difficult, particularly in the context of investment decisions, to establish a sufficiently comprehensive and reliable decision basis for top management. In such situations, the approach commonly described as "start small, scale fast" has proven to be effective. This means that a limited pilot or preliminary software project may be required to justify a subsequent, larger investment in software-based solutions.

The measures presented for implementing predictive maintenance indicate that the greatest value is generated primarily in the preparatory and follow-up phases of actual maintenance execution. Figure 13 provides a structured overview illustrating this relationship. While insights gained from predictive maintenance can be transferred into digital assistance systems to support maintenance execution, human labor remains essential for carrying out maintenance activities. Digital assistance systems, such as smart glasses with augmented reality functionality for real-time guidance, can support technicians during execution. However, the sensory and motor skills required for maintenance tasks cannot be automated in the short to medium term.

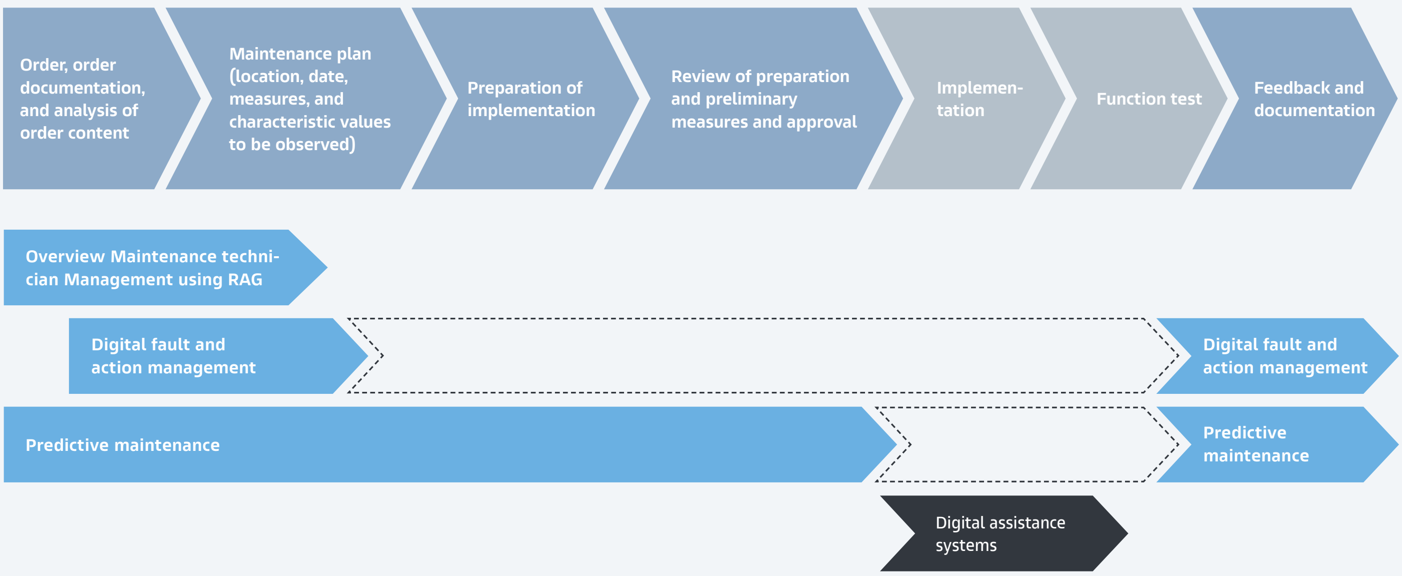

Figure 13: Overview of measures for implementing predictive maintenance compared with the repair phase

In summary, software support primarily reduces the administrative effort associated with maintenance activities.
As a result, KPIs related to personnel costs per maintenance case form a significant basis for ROI calculations. At the same time, improved maintenance planning enables a substantial reduction in spare parts inventory. Lower inventory levels reduce capital tie-up, increase capital turnover, and ultimately lead to higher returns on capital.
Indirect cost effects, such as production losses, are more difficult to quantify, but they can play an important role in parallel risk assessments. Additional indirect effects, including the benefits of structured knowledge management, should also be considered. Company-specific and equipment-specific knowledge can be prepared on demand by software systems and made available to maintenance technicians immediately before execution. As a result, technicians receive more targeted information about the machines and upcoming tasks and can work more effectively.

Effects of predictive maintenance supported by IIoT and AI[2]

18 – 35%

Reduction of operating costs

Increase in equipment availability

2 Source: Majeed, H., Iftikhar, T. (2026): Intelligent Manufacturing in Industry 6.0. A Climate Resilience Approach. Springer.

5. Implementing predictive maintenance using the industrial data platform concept

The Industrial Data Platform concept provides substantial support for predictive maintenance by serving as a central foundation for capturing, integrating, storing, and analyzing large volumes of heterogeneous machine data as well as data from other software systems used in production. In many manufacturing companies, predictive maintenance initiatives have failed due to inconsistent data, fragmented data sources, or limited accessibility across different layers of the production environment. A central data platform, operated either in the cloud or on-premises servers, addresses these challenges by aggregating data from machines, sensors, ERP systems, and MES systems in a unified manner and providing structured and consistent access to this data.

A further advantage of a central data platform is the ability to exchange data securely with external maintenance service providers or suppliers, for example, through standardized interfaces. In many cases, operators are unable to develop all the required analysis models for machines on their own. Partnership-based collaboration with suppliers and external specialists, such as engineering service providers, can therefore be significantly simplified through governed and standardized data exchange. External experts can access secured live data and digital twins to develop and refine analysis models. As a result, service deployments can be planned more precisely, and spare parts demand can be identified at an early stage.

A technical discussion of the concept is provided in the white paper "Industrial Data Platform".

Overall, the Industrial Data Platform enables the development and implementation of scalable predictive maintenance solutions. It provides the foundation for reducing the risk of unplanned machine downtime and optimizing maintenance intervals. As a result, the efficiency and reliability of production processes increase, and overall equipment effectiveness (OEE) improves significantly. Experience from the semiconductor industry demonstrates a clear additional benefit. Companies that use AI-supported fault detection systems have achieved yield increases of up to 20 percent, primarily due to reductions in waste material and process-related losses.
This illustrates that data-driven fault detection and predictive maintenance not only improve equipment availability but also have a direct and measurable positive impact on yield and material efficiency.

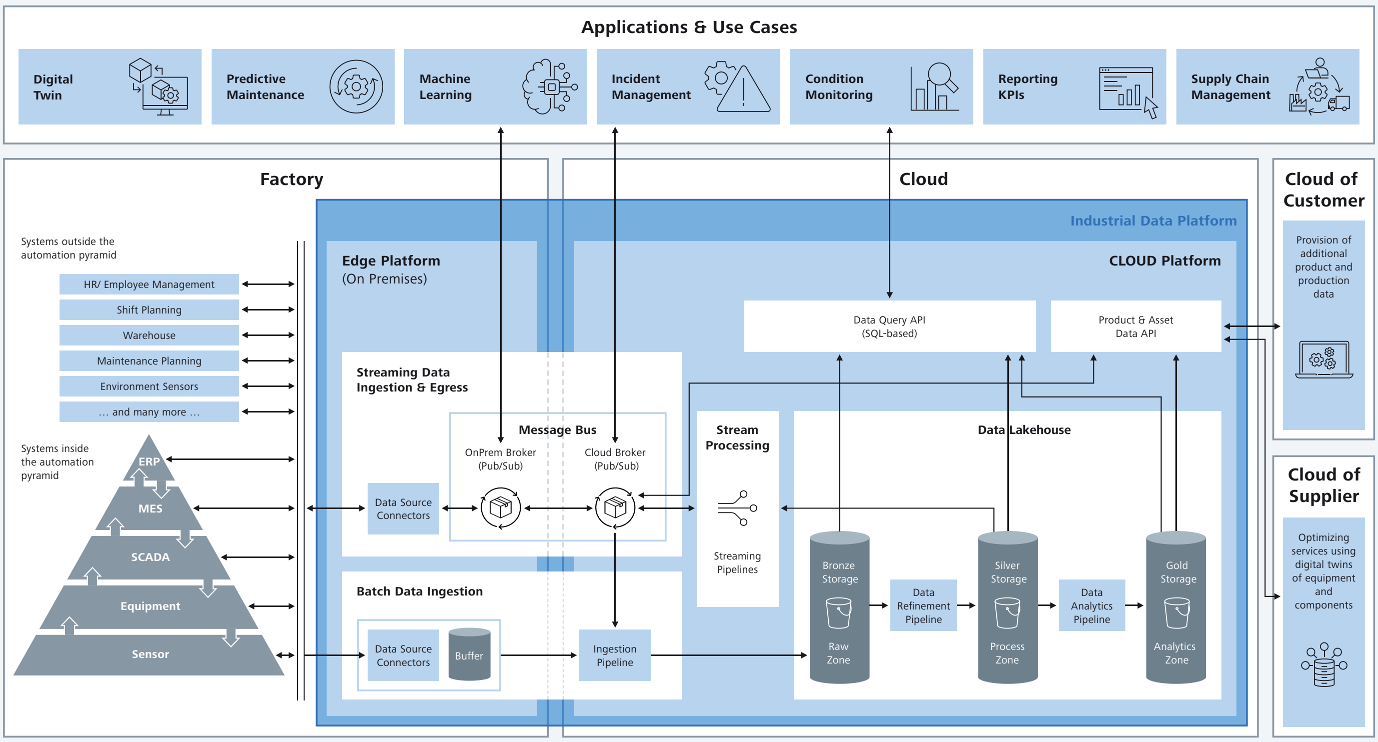

Figure 14: Industrial Data Platform concept, data-flow view

6. Conclusion

In the pharmaceutical industry, unplanned equipment downtime can result in substantial shortage-related costs, while cyclic maintenance can place a significant burden on maintenance cost centers. Predictive maintenance, therefore, represents a decisive solution approach. This white paper derives an implementation path by first discussing the theoretical foundations of maintenance. It then identifies process objectives and outlines a step-by-step implementation strategy. The approach starts with the digitalization of maintenance information and digital fault and action management and progresses toward full predictive maintenance based on a central data platform. The primary objective is to replace reactive or rigid time-based maintenance strategies with a dynamic, condition-oriented approach.
This reduces unplanned downtime, optimizes mainte-nance cycles, and lowers both direct and indirect maintenance costs. Over the long term, a significant increase in overall equipment effectiveness (OEE) is targeted.
The strategic implementation of a central data platform enables the integration of heterogeneous data sources and provides the foundation for building and analyzing data models. Suitable KPIs are proposed to support the required ROI assessment. The Industrial Data Platform concept forms the basis for scalable predictive maintenance solutions, which are essential for maintaining competitiveness in the high-technology industry.

As a digitalization partner with deep expertise in highly automated production environments, ZEISS Digital Innovation supports companies in translating this target state into operational reality at the intersection of domain expertise, data, and technology. The objective is to secure targeted investments and to design digital solutions that deliver measurable value in day-to-day operations.

Author

Marco Grafe
Solution Specialist for Smart Manufacturing at ZEISS Digital Innovation
marco.grafe@zeiss.com

Get in touch to exchange about your ideas and how we can help.

Carl Zeiss Digital Innovation GmbH
Fritz-Foerster-Platz 2
01069 Dresden
Germany

Phone: +49 351 49701 – 500
contact.digitalinnovation.de@zeiss.com
zeiss.com/digital-innovation