International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246.IJETCSIT-V7I2P101
Eureka Vision Publication | Volume 7, Issue 2, 1-11, 2026

Original Article

Predictive Maintenance Analytics Using Dynamics 365 Supply Chain Management and Microsoft Fabric

Manish Sonthalia
Independent Researcher, USA.

Received On: 11/02/2026 Revised On: 17/03/2026 Accepted On: 25/03/2026 Published On: 02/04/2026

Abstract - *Unplanned equipment downtime continues to rank among the most costly disruptions in modern manufacturing and supply chain operations. This whitepaper examines how the integration of Microsoft Dynamics 365 Supply Chain Management (D365 SCM) and Microsoft Fabric delivers a unified, end-to- end platform for predictive maintenance analytics. We explore the core challenges that drive organizations away from reactive and schedule-based maintenance strategies toward data-driven, condition-based approaches including the high cost of emergency repairs, wasted resources from over-servicing healthy equipment, and the growing complexity of managing geographically distributed asset fleets. The paper details how D365 SCM's Asset Management module and Sensor Data Intelligence feature provide the operational backbone for asset tracking, work order management, and IoT integration, while Microsoft Fabric's unified analytics architecture anchored by One Lake, Synapse Data Science, Real-Time Intelligence, and Power Bide livers the analytical engine for ingesting sensor telemetry, engineering features, training machine learning models, and surfacing actionable predictions. Key outcomes documented across industry implementations include 35–50% reductions in unplanned downtime, 25–35% maintenance cost savings, and 20–40% extensions in asset operational life. The paper concludes with a phased implementation methodology, best practices for data governance and model lifecycle management, and a framework for quantifying return on investment.*

Keywords - *Predictive Maintenance, Dynamics 365 Supply Chain Management, Microsoft Fabric, Internet of Things (IoT), Machine Learning, Asset Management, Predictive Analytics, Remaining Useful Life (RUL).*

1. Introduction to Predictive Maintenance

Manufacturing and supply chain operations depend on physical assets compressors, turbines, conveyor belts, CNC machines, packaging lines and every one of those assets will eventually degrade. The question is not whether maintenance will be needed, but when and how it should be performed. For decades, organizations have toggled between two imperfect strategies. Reactive maintenance waits for equipment to break, then scrambles to fix ita model that maximizes unplanned downtime and emergency-repair costs. Preventive maintenance applies service at fixed intervals, which eliminates some surprises but introduces its own inefficiency: healthy components are replaced too early, and faults that develop between scheduled checks still cause unexpected failures.

Predictive maintenance (PdM) offers a third way. By attaching IoT sensors to equipment and streaming their readings vibration, temperature, pressure, acoustic emissions, electrical current to a cloud analytics platform, organizations can monitor asset condition around the clock. Machine learning algorithms learn what "normal" looks like for each machine and flag subtle deviations that signal the early stages of wear or fault development. The output is not just an alarm; it is an estimated time to failure, often expressed as Remaining Useful Life (RUL), that lets planners schedule repair during a convenient window rather than endure a disruptive breakdown.

The appeal is straightforward. Maintenance happens only when the data says it is needed, so organizations avoid both the cost of doing nothing and the waste of doing too much. Parts are ordered in advance, technicians are allocated efficiently, and production loses the minimum possible time. What makes this practical today rather than theoretical is the convergence of affordable IoT hardware, scalable cloud platforms, and mature AI/ML tooling. Microsoft's ecosystem, specifically Dynamics 365 Supply Chain Management for operational execution and Microsoft Fabric for unified analytics, provides a ready-made foundation for this transformation.

2. The Business Case for Predictive Maintenance

Building a predictive maintenance program requires real investment sensors, connectivity, platform licenses, data science talentso any proposal needs a rigorous business case. Fortunately, the data is persuasive. Industry benchmarks consistently show that PdM reduces unplanned downtime by 35–50%, cuts overall maintenance costs by 25–35% compared with preventive strategies, and extends equipment life by 20–40%. McKinsey has estimated that predictive maintenance could save companies as much as $630 billion globally by eliminating unnecessary servicing alone.

The savings come from several directions at once. First, by replacing calendar-based schedules with condition-based triggers, PdM eliminates over-maintenance no more pulling a perfectly functional bearing off a motor because the manual says to inspect it every 2,000 hours. Second, catching faults early prevents the cascading damage that turns a $500 repair into a $50,000 rebuild. Third, because repairs are planned rather than emergency-driven, labor can be scheduled at standard rates instead of overtime, and spare parts can be procured through normal channels rather than expedited shipping.

Beyond the direct financial gains, PdM improves safety by catching hazardous conditionsexcessive vibration, overheating, abnormal pressurebefore they endanger personnel. It also supports sustainability goals: well-maintained equipment runs more efficiently, consuming less energy and producing less scrap. And in highly regulated sectors such as pharmaceuticals and food processing, continuous condition monitoring generates an audit trail that simplifies compliance reporting. For most organizations, the return on investment materializes within 6 to 18 months of initial deployment, with benefits compounding as the program scales to cover additional asset classes.

3. Dynamics 365 Supply Chain Management Capabilities

Dynamics 365 Supply Chain Management sits at the operational heart of the predictive maintenance solution. Its Asset Management module provides a single system of record for every physical asset an organization owns or operatesfrom serial-numbered motors and pumps to entire production lines. Each asset carries a rich profile: location hierarchy, technical specifications, maintenance history, warranty information, and associated spare-parts bills of material. This structured data is essential context for any machine learning model, because predicting when a compressor will fail requires knowing not just its current sensor readings but also its age, maintenance history, and operating environment.

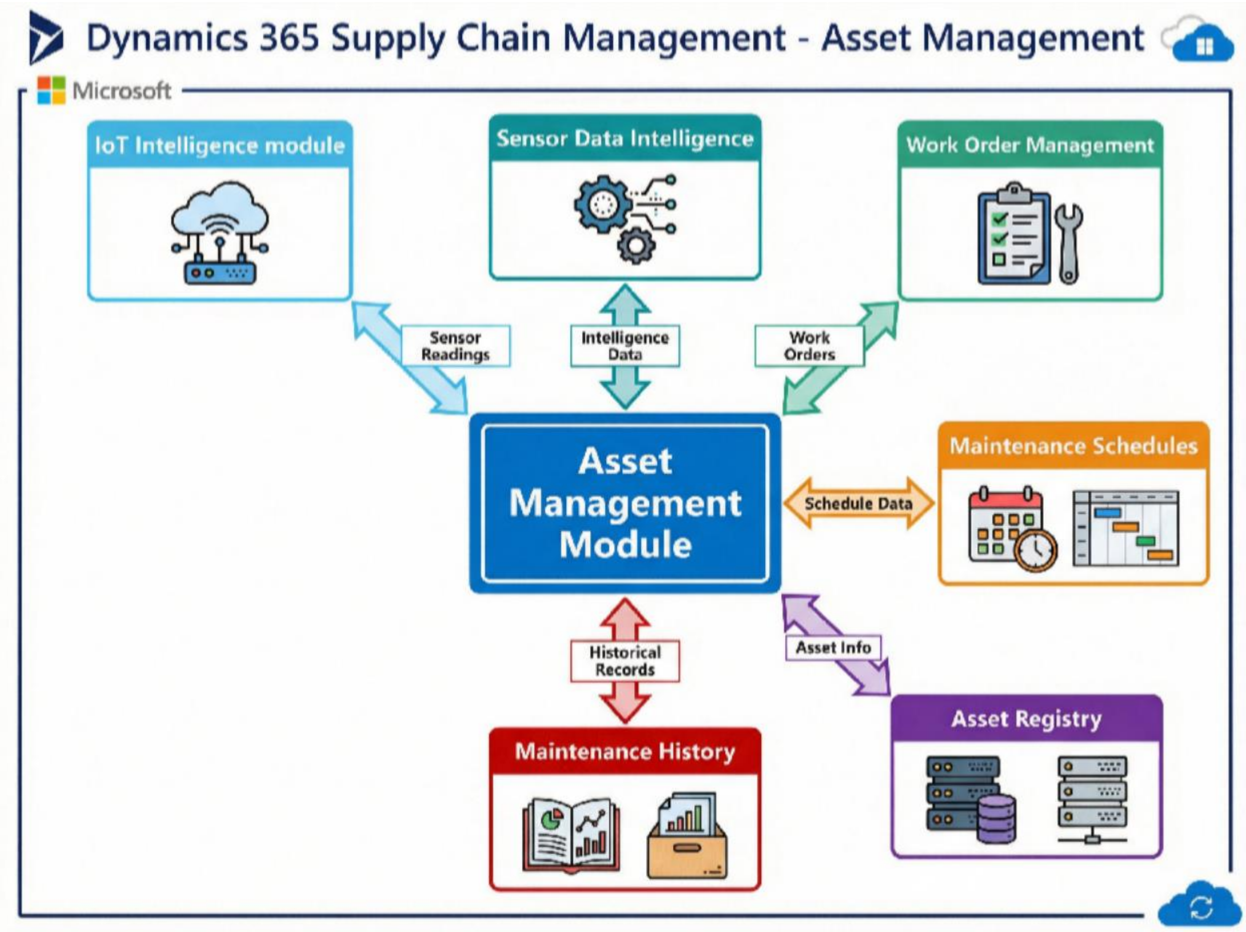

Figure 1. D365 Scm Asset Management Architecture

Work order management is the module's transactional backbone. When a maintenance plan fires whether triggered by a calendar schedule, a usage counter, or a predictive alert from FabricD365 SCM generates a work order that captures the required tasks, skill requirements, estimated duration, and needed parts. Technicians receive assignments on mobile devices, can view checklists and asset documentation on-site, and log actuals (time, materials, findings) directly in the system. All costs flow through to project accounting and general ledger, giving finance teams full visibility into maintenance spend by asset, location, or failure type.

What truly differentiates D365 SCM in the predictive maintenance context is Sensor Data Intelligence (SDI). This feature deploys Azure IoT Hub and Azure Stream Analytics resources directly into the customer's own Azure subscription, giving IT teams full control over networking, security, and scaling. SDI supports multiple business scenarios out of the box: Asset Maintenance monitors sensor parameters to trigger condition-based work orders; Asset Downtime tracks machine availability for OEE calculations; Machine Status alerts production planners to equipment outages; Production Delays compares actual and planned cycle times; and Product Quality flags batches that drift outside specification. Together, these scenarios turn D365 SCM from a maintenance management system into a real-time, IoT-aware operational platform.

4. Microsoft Fabric for Predictive Analytics

While D365 SCM handles the operational sidetracking assets, dispatching technicians, recording coststhe heavy analytical lifting happens in Microsoft Fabric. Fabric is Microsoft's unified, AI-powered analytics platform, delivered as a single SaaS offering that brings together capabilities previously spread across Azure Data Factory, Azure Synapse Analytics, and Power BI. The unification matters because predictive maintenance is inherently a multi-discipline problem: it requires data engineers to build ingestion pipelines, data scientists to train models, real-time analysts to monitor streaming telemetry, and business users to interpret dashboards. Fabric gives each of those personas a tailored experience while keeping everyone on the same platform and the same copy of data.

At Fabric's core is OneLake, a single, organization-wide data lake built on Azure Data Lake Storage Gen2 and the open Delta Lake Parquet format. Every Fabric workload reads from and writes to OneLake, which means that a Spark job in Synapse Data Engineering, a machine learning experiment in Synapse Data Science, and a Power BI report all operate on the same physical datano copying, no synchronization scripts, no version-mismatch headaches. OneLake also supports Shortcuts, virtual pointers that let you reference data in external storage accounts (Azure, AWS S3, Google Cloud) without moving it, and forthcoming Mirroring capabilities that will replicate D365 operational data into OneLake in near-real time.

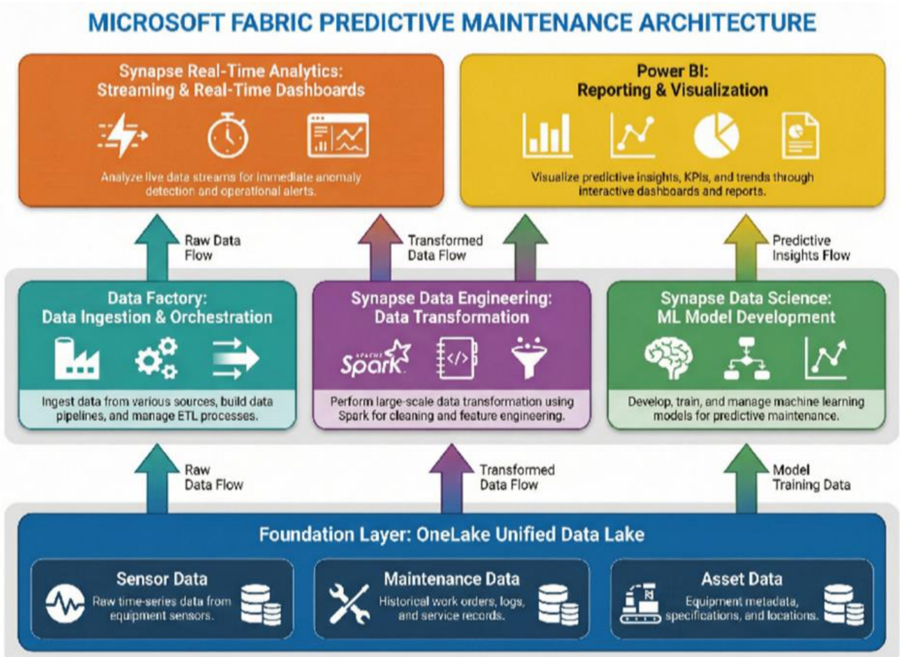

Figure 2. Microsoft Fabric Analytics Architecture for Predictive Maintenance

Fabric's workloads map neatly to the predictive maintenance pipeline. Data Factory handles batch ingestion and transformation with over 200 native connectors, low-code Dataflow Gen2 transforms, and orchestration pipelines that support change data capture. Synapse Data Engineering provides Apache Spark for large-scale feature engineering on years of historical sensor data. Synapse Data Science offers notebooks, MLflow experiment tracking, and the PREDICT function for batch scoring. Real-Time Intelligence captures streaming IoT data through Eventstreams and Eventhouses, with Kusto Query Language (KQL) powering sub-second anomaly detection. Power BI visualizes predictions through Direct Lake modeconnecting directly to OneLake Delta tables without data importand Data Activator monitors metrics continuously and triggers automated actions when thresholds are breached.

5. Predictive Maintenance Framework

A successful predictive maintenance initiative is not a single technology deployment but a coordinated framework that spans people, process, and technology. At the highest level, the framework connects three domains: the physical domain (assets, sensors, edge gateways), the data and analytics domain (ingestion, storage, modeling, visualization), and the operational domain (work order execution, spare parts management, continuous improvement). The diagram below illustrates how these domains interact in a closed-loop system where sensor signals ultimately drive maintenance actions, and the outcomes of those actions feed back into the models to improve future predictions.

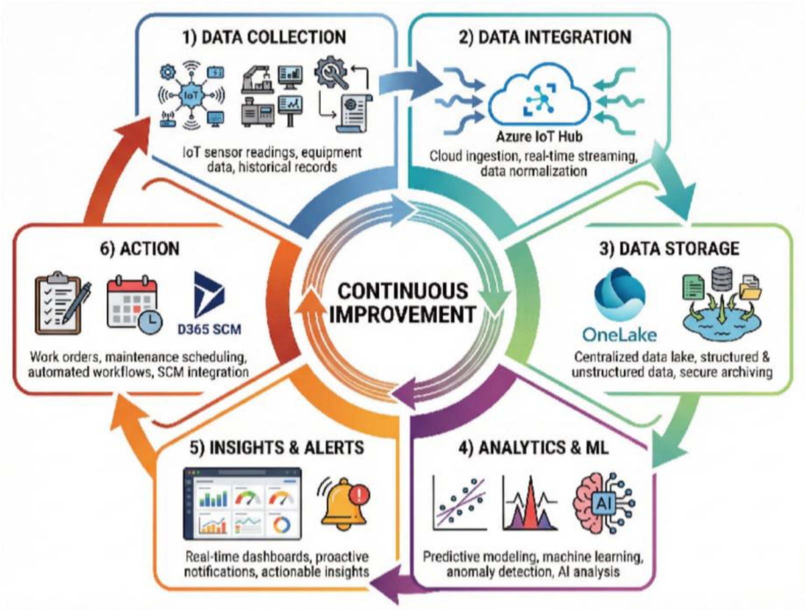

Figure 3. Predictive Maintenance Framework Overview

Within the physical domain, the priority is instrumenting the right assets with the right sensors. Not every machine warrants full IoT coverage; a criticality assessment should identify the assets whose failure carries the greatest cost in downtime, safety risk, or downstream impact. For those assets, sensor selection depends on the failure modes to be detectedaccelerometers for bearing wear, thermocouples for overheating, current sensors for motor degradation, acoustic emission sensors for leak detection. Edge gateways aggregate and pre-filter the high-frequency data streams before transmitting them to the cloud, reducing bandwidth costs and latency.

The data and analytics domain is where Microsoft Fabric plays its central role. Raw sensor telemetry is ingested through Eventstreams for real-time processing and through Data Factory for batch loads. Feature engineering transforms the raw signals into meaningful predictorsrolling averages, spectral features, rate-of-change metricsthat the machine learning models consume. Trained models produce two primary outputs: anomaly scores for immediate alerting and RUL estimates for medium-term planning. Both feed into Power BI dashboards and Data Activator triggers that connect to the operational domain via D365 SCM, ensuring that every prediction reaches the people and systems that can act on it.

6. End-To-End Architecture and Data Flow

Understanding the complete data journeyfrom a vibration reading on a factory-floor motor to a work order on a technician's tabletis essential for architects planning the solution. The architecture follows four distinct phases: ingestion, processing and storage, analysis and modeling, and action.

Ingestion begins at the edge. IoT sensors capture telemetry at frequencies ranging from once per second for temperature to thousands of samples per second for vibration. Edge devices perform initial filtering and aggregation, then transmit the data to Azure IoT Hub over secure MQTT or AMQP connections. Within Fabric, an Eventstream subscribes to the IoT Hub and routes the streaming data to an Eventhouse for real-time querying. In parallel, Data Factory pipelines pull contextual data from D365 SCMasset registries, maintenance history, spare-parts catalogsand land it in a Lakehouse in OneLake. The co-location of sensor data and business context in the same data lake is what makes cross- domain analytics possible.

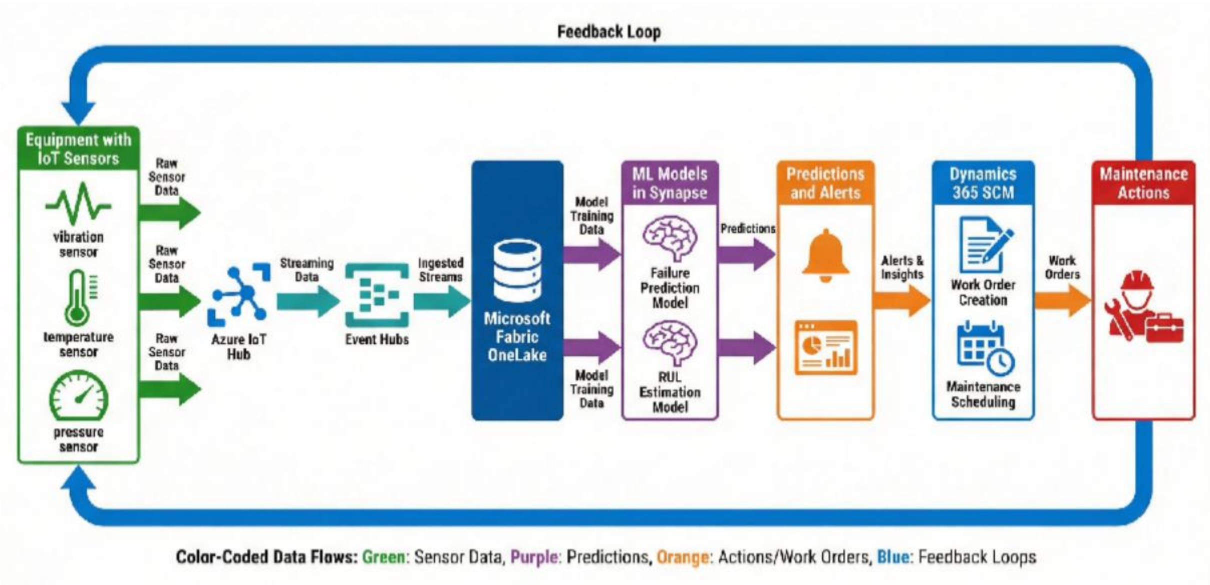

Figure 4. End-To-End Data Flow Architecture

Processing and storage happen on two parallel tracks. The real-time track uses KQL queries running against the Eventhouse tocompute sliding-window aggregates, detect threshold breaches, and apply built-in anomaly-detection functions to streaming telemetry. The batch track uses Spark in Synapse Data Engineering to run nightly or hourly jobs that join historical sensor data with maintenance records, compute advanced features (e.g., frequency-domain transforms, degradation slopes), and write the results to curated Delta tables in the Lakehouse.

Analysis and modeling consume the curated data. Data scientists in Synapse Data Science train classification models (healthy vs. degrading vs. critical) and regression models (days until failure) using libraries such as XGBoost, LightGBM, and PyTorch. MLflow tracks every experiment, and the best- performing model is registered in the model registry. Batch scoring runs daily via the Spark PREDICT function, writing updated RUL estimates and health scores to a predictions table in the Lakehouse.

The action phase closes the loop. Power BI reports built in Direct Lake mode visualize asset health across the fleetheat maps of risk, trend lines of degradation, drill-downs to individual sensor channels. Data Activator monitors the predictions table and the real-time Eventstream; when an asset's predicted RUL drops below a configurable threshold (say, 30 days), it fires a Power Automate flow that creates a priority work order in D365 SCM, pre-populated with the affected asset, recommended tasks, and diagnostic context from the model. The technician sees the work order on a mobile device, performs the repair, logs actuals, and the updated maintenance record flows back into Fabric to retrain the model completing the feedback loop.

7. Machine Learning Models and Algorithms

The predictive power of the solution rests on the machine learning models that translate raw sensor signals into actionable forecasts. In practice, predictive maintenance draws on three broad classes of algorithms, each suited to a different question.

Classification models answer the question "What state is this asset in right now?" Gradient-boosted tree ensemblesXGBoost and LightGBM in particularare popular choices because they handle tabular data well, tolerate missing values, and train quickly. A typical model takes as input a feature vector composed of rolling statistics (mean, standard deviation, kurtosis) computed over recent sensor windows, plus contextual features such as asset age and last-service date, and outputs a probability distribution across health states (e.g., normal, watch, warning, critical). The thresholds that map probabilities to operational actions are tuned in consultation with maintenance engineers who understand the cost trade-offs between false positives (unnecessary work orders) and false negatives (missed failures).

Regression models tackle the question "How much useful life remains?" Estimating Remaining Useful Life (RUL) is fundamentally a time-to-event prediction problem. For assets with smooth degradation curvesbearings, for example, whose vibration amplitude rises steadily before failureclassical survival analysis or simple regression on degradation features can be effective. For assets with more complex or irregular degradation patterns, deep learning architectures such as Long Short-Term Memory (LSTM) networks and Temporal Convolutional Networks (TCN) capture sequential dependencies in the sensor time series. These models are trained on run-to-failure datasets where the true RUL is known, and they learn to map a window of recent sensor readings to a predicted number of remaining operational cycles or hours.

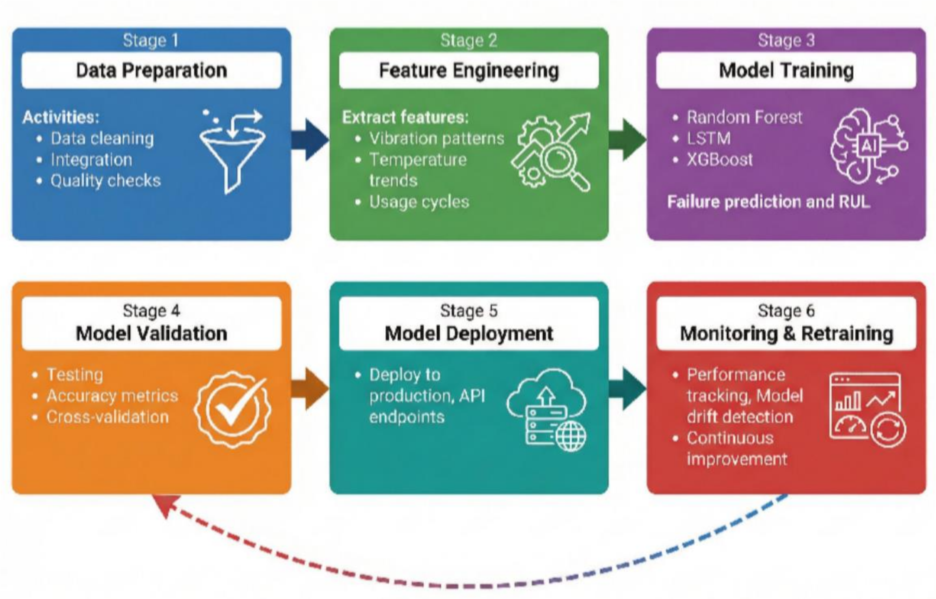

Figure 5. Machine Learning Pipeline for Predictive Maintenance

Anomaly detection models address a different but complementary need: "Is something unusual happening right now?" Autoencodersneural networks trained to reconstruct normal operating data are particularly effective here. When the reconstruction error spikes, it signals that the current sensor pattern deviates from anything the model learned during training, flagging a potential new or rare fault. Isolation Forests and One-Class SVMs offer lighter-weight alternatives that work well when labeled failure data is scarce. In the Fabric architecture, anomaly detection typically runs on the real-time track (via KQL functions or lightweight models deployed at the edge), while classification and RUL models run as batch jobs in Synapse Data Science.

Regardless of algorithm choice, model lifecycle management is critical. Fabric's native MLflow integration tracks every training runhyperparameters, evaluation metrics (precision, recall, RMSE), training data versionensuring full reproducibility. Models are registered, versioned, and promoted through staging environments before reaching production. Automated retraining pipelines, triggered by data drift detection or on a fixed schedule, ensure that models stay current as equipment ages and operating conditions evolve.

8. Use Cases and Applications

Predictive maintenance is not a one-size-fits-all proposition; its value varies by industry and asset type. The following scenarios illustrate how the D365 SCM and Fabric platform adapts to diverse operational contexts.

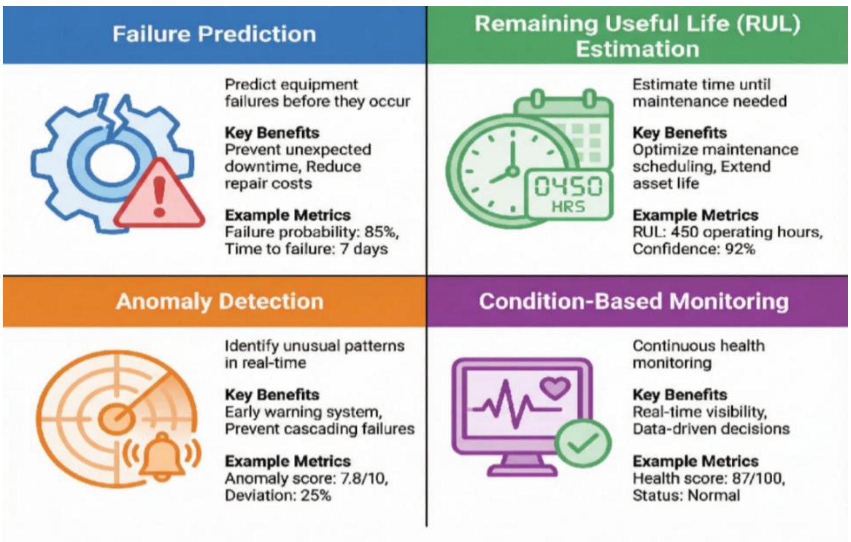

Figure 6. Predictive Maintenance Use Cases across Industries

In discrete manufacturing, CNC machines, robotic welding arms, and stamping presses are prime candidates.

Vibration andcurrent sensors on spindle bearings feed models that predict tool wear, enabling just-in-time tool changes that prevent scrap without wasting tool life. An automotive parts manufacturer, for example, might monitor hundreds of CNC spindles across multiple plants. Fabric ingests the telemetry, Synapse Data Science trains a spindle-specific RUL model, and when a spindle's predicted life drops below the next shift duration, Data Activator creates a work order in D365 SCM so the tool change happens during the shift break.In process industriesoil and gas, chemicals, food and beveragecontinuous processes run on pumps, compressors, and heat exchangers where even brief outages are enormously expensive. Temperature and pressure sensors on a refinery's critical compressors, analyzed by LSTM models in Fabric, can detect subtle degradation in valve seals weeks before a leak occurs. The D365 SCM work order includes not just the repair task but also the specific spare parts and hazardous-area work permits required, streamlining compliance.

Energy and utilities offer another compelling arena. Wind farm operators use vibration, oil-debris, and SCADA data to predict gearbox and generator-bearing failures across turbine fleets. Because turbine nacelles are difficult to access and require specialized cranes, predicting failures far enough in advance to schedule a maintenance vessel and crew yields massive logistics savings. Power BI dashboards display fleet-wide health maps, letting operations managers prioritize which turbines to service during the next weather window.

In logistics and warehousing, conveyor systems, sortation equipment, and automated guided vehicles (AGVs) are critical to throughput. Current sensors on conveyor motors detect belt misalignment or roller degradation, and acoustic sensors on AGV drive wheels flag bearing wear. Because these environments run around the clock during peak seasons, the ability to schedule preventive swaps during brief maintenance windows is essential to meeting delivery commitments.

9. Iot Integration and Real-Time Monitoring

The sensor layer is the nervous system of a predictive maintenance solutionwithout reliable, high- quality data flowing from the physical world, even the most sophisticated models are useless. A well- designed IoT integration strategy addresses three concerns: sensor selection, connectivity architecture, and real-time processing.

Sensor selection starts with a Failure Mode and Effects Analysis (FMEA) for each critical asset. The FMEA identifies the most likely and most consequential failure modes, and for each mode, the physical parameter that best indicates early degradation. Vibration accelerometers are indispensable for rotating equipment (motors, pumps, fans), while infrared temperature sensors detect thermal anomalies in electrical cabinets and bearings. Pressure transducers monitor hydraulic systems and sealed vessels, and ultrasonic microphones pick up the high-frequency emissions characteristic of early-stage leaks and electrical discharge. In many cases, a combination of sensor types is needed to distinguish between failure modes that produce similar symptoms in a single parameter.

Connectivity architecture determines how sensor data reaches the cloud. For brownfield plants, wireless protocols like LoRaWAN, Bluetooth Low Energy, and industrial Wi-Fi minimize the cost of retrofitting cabling. Edge gateways aggregate data from multiple sensors, apply initial filtering (e.g., discarding readings during planned shutdowns), and batch-transmit to Azure IoT Hub over TLS-encrypted connections. For latency-sensitive use casessuch as detecting a sudden pressure spike that requires an immediate machine stopedge computing nodes running Azure IoT Edge can execute lightweight inference models locally, triggering safety interlocks in milliseconds before the data even reaches the cloud.

Once data arrives in Microsoft Fabric, the Real-Time Intelligence workload takes over. Eventstreams ingest the IoT Hub data and route it to an Eventhouse, where KQL queries run continuously to compute sliding-window statistics, detect threshold breaches, and score anomaly-detection models. The Eventhouse retains high-resolution data for a configurable retention period (commonly 30–90 days) before aggregated summaries are archived to the Lakehouse for long-term historical analysis. This tiered storage approach balances query performance with cost efficiency and ensures that both real-time dashboards and historical model-training datasets are served from the same platform.

10. Implementation Methodology

Deploying a predictive maintenance solution is a multi-phase journey, not a big-bang project. A phased approach manages risk, builds organizational confidence, and delivers incremental value at each stage.

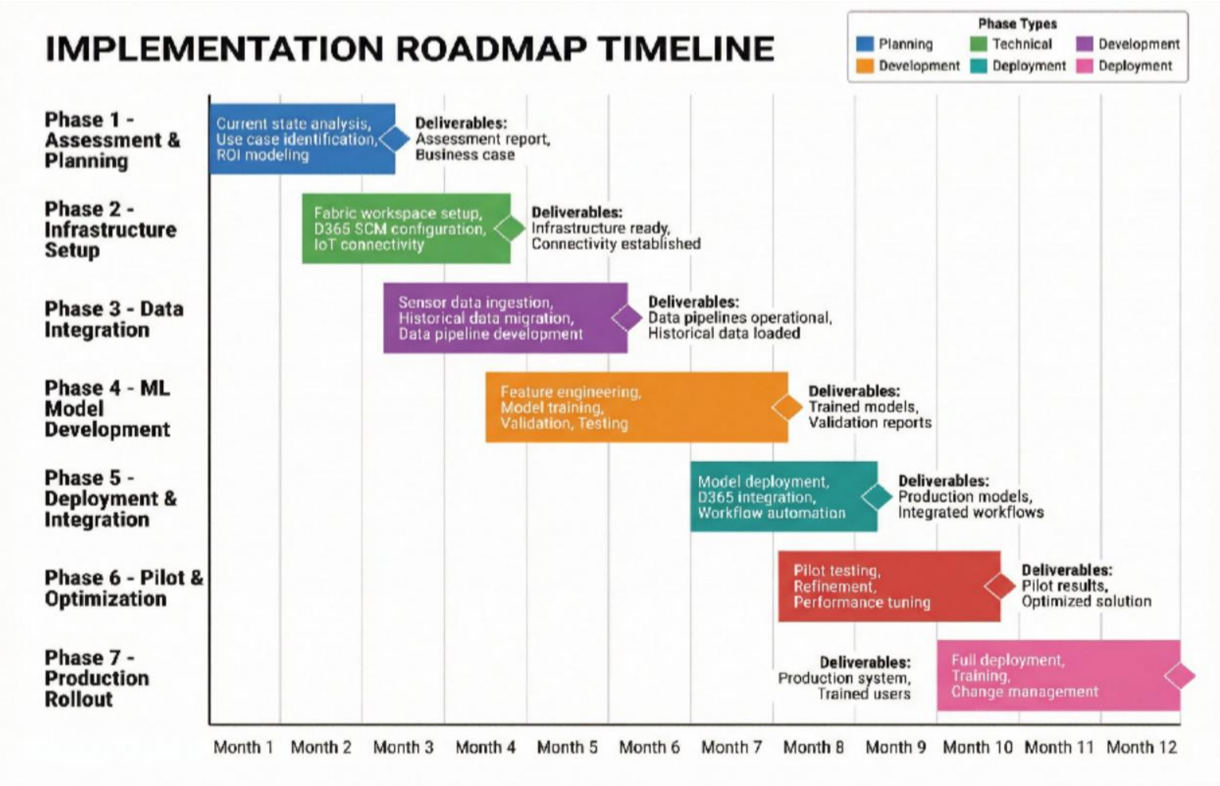

Figure 7. Implementation Roadmap

- Phase 1: Assessment and Planning (4–6 weeks). The engagement begins with a criticality assessment that ranks assets by the business impact of their failure. A cross-functional teammaintenance engineers, production managers, IT architects, data scientistsselects 3–5 pilot assets that are both high-impact and well-instrumented (or easy to instrument). The team documents existing maintenance processes, identifies data sources, and defines success metrics: target reductions in unplanned downtime, maintenance cost per asset, mean time between failures. A high-level solution architecture is drafted, mapping data flows from sensors through Fabric to D365 SCM.
- Phase 2: Pilot Implementation (8–12 weeks). Sensors are installed on the pilot assets, and the data pipeline is built end to endfrom IoT Hub ingestion through Fabric Eventstreams and Lakehouses to initial Power BI dashboards. Data engineers develop feature-engineering notebooks, and data scientists begin exploratory modeling using the accumulating historical data. During this phase, the team typically runs the predictive system in "shadow mode," where it generates alerts alongside the existing maintenance process without directly triggering work orders. This allows the models to be validated against real outcomes before they influence operations.
- Phase 3: Operationalization (6–8 weeks). Once the pilot models demonstrate acceptable accuracy (precision, recall, and RUL error within agreed thresholds), the system transitions to live operation. Data Activator triggers are connected to D365 SCM to auto-generate work orders, and maintenance planners begin scheduling based on RUL predictions rather than fixed intervals. Change management is critical in this phase: technicians need training on the new mobile workflows, planners need to trust the model outputs, and managers need dashboards that surface the right KPIs. A feedback loop is established so that maintenance outcomes (confirmed failures, false alarms, parts replaced) are captured and fed back into model retraining.
- Phase 4: Scaling and Optimization (ongoing). With the pilot proven, the program expands to additional asset classes and facilities. Each new asset type may require new sensors, new features, and new or fine-tuned models, but the underlying Fabric platform and D365 integration patterns are reused. Over time, the growing volume of failure data improves model accuracy, and the organization matures its data governance, model-ops practices, and continuous-improvement culture. Advanced capabilities such as prescriptive maintenance (recommending the optimal repair action, not just that repair is needed) and digital twinscan be layered onto the established foundation.

11. ROI and Business Benefits

Quantifying the return on investment from predictive maintenance requires mapping the solution's technical outputs to financial outcomes. The primary value drivers fall into four categories: downtime reduction, maintenance cost optimization, asset life extension, and secondary operational improvements.

Downtime reduction is typically the largest single contributor to ROI. When a critical production line goes down unexpectedly, the cost is not just the repair itself but lost production, spoiled work-in- progress, expedited shipping to compensate for delays, and potential contractual penalties. Industry data suggests that unplanned downtime costs manufacturers an average of $260,000 per hour in some sectors. A 35–50% reduction in unplanned events can therefore translate to millions of dollars in annual savings for a large facility.

Figure 8. ROI and Benefits Dashboard

Maintenance cost optimization covers both labor and materials. By eliminating unnecessary preventive tasks and shifting emergency repairs to planned maintenance, organizations see 25–35% reductions in total maintenance expenditure. Labor savings come from scheduling work during regular shifts instead of overtime, batching tasks for efficiency, and reducing the diagnostic time that technicians spend hunting for faults (since the model provides a probable diagnosis). Material savings arise from ordering parts through standard procurement channels and from extending the service life of components that would have been replaced prematurely under a preventive schedule.

Asset life extension compounds these savings over the long term. Equipment that is maintained at optimal intervalsneither too early nor too latesuffers less cumulative damage and operates within its design parameters for longer. A 20–40% extension in operational life defers capital expenditure on replacements, freeing budget for growth investments. Secondary benefits include improved safety (fewer hazardous failure events), higher product quality (equipment operating within spec produces fewer defects), and enhanced sustainability metrics (lower energy consumption and waste). When all these factors are aggregated, organizations typically achieve full payback on their PdM investment within 6–18 months, with compounding returns as the program matures and scales.

12. Best Practices and Recommendations

Having examined the technology, architecture, and business case, the following recommendations distill the lessons learned from successful implementations.

Prioritize data quality over model complexity. A simple gradient-boosted model trained on clean, well- engineered features will outperform a deep neural network fed noisy, poorly labeled data every time. Invest in robust data validation at ingestion, clear labeling of historical failure events, and systematic feature engineering before reaching for exotic algorithms. Fabric's Data Factory and Data Engineering workloads provide the tools, but discipline in using them is what makes the difference.

Adopt a hybrid maintenance strategy. Predictive maintenance is most valuable for high-criticality, high- cost assets where the ROI justifies the instrumentation and modeling investment. For low-criticality equipmentlight fixtures, office HVAC units, non-bottleneck conveyor sectionssimple preventive schedules or even run-to-failure strategies may remain the most cost-effective approach. D365 SCM's Asset Management module supports all three strategies within a single system, allowing organizations to apply the right approach to each asset class.

Build cross-functional teams. Predictive maintenance sits at the intersection of operations technology (OT), information technology (IT), and data science. Projects that silo these disciplines tend to produce technically impressive models that nobody uses, or operationally sound processes that lack analytical rigor. The most successful programs embed data scientists within maintenance teams, give maintenance engineers a voice in feature selection, and ensure IT architects design for production-grade reliability from the start.

Plan for model lifecycle management from day one. Models degrade as equipment ages, operating conditions shift, and new failure modes emerge. Fabric's MLflow integration supports experiment tracking and model versioning, but organizations also need automated pipelines that monitor model performance in production, detect data drift, and trigger retraining when accuracy falls below threshold. Treat ML models as living assets that require ongoing carenot unlike the physical assets they monitor.

Invest in change management. The shift from calendar-based to prediction-based maintenance changes how planners, technicians, and managers work. Transparent communication about why the change is happening, hands-on training with the new tools (mobile work orders, Power BI dashboards), and early wins that demonstrate tangible benefits are all essential to sustained adoption. Leadership sponsorship ensures that the program survives the inevitable setbacks and resource competitions of the first year.

13. Conclusion

The convergence of IoT, cloud analytics, and enterprise resource planning has created a genuine inflection point for industrial maintenance. Organizations that continue to rely solely on reactive repairs or rigid preventive schedules will increasingly find themselves at a competitive disadvantagespending more on maintenance, suffering more downtime, and replacing equipment sooner than necessary. The integrated platform of Dynamics 365 Supply Chain Management and Microsoft Fabric offers a clear, proven path to a better model.

D365 SCM provides the operational foundation: a structured system of record for assets, maintenance plans, and work orders, enriched by Sensor Data Intelligence's real-time IoT bridge. Microsoft Fabric supplies the analytical engine: a unified platform where data engineers, data scientists, and business analysts collaborate on a single copy of datafrom raw sensor telemetry to curated predictions without the friction of stitching together disparate services. The closed-loop architecture ensures that every predictive insight translates into a concrete action: a work order dispatched, a spare part ordered, a technician briefed with diagnostic context.

The benefits documented across early adopters are compelling and consistent: 35–50% reductions in unplanned downtime, 25–35% lower maintenance costs, 20–40% longer asset lifespans, and measurable improvements in safety and product quality. Achieving these outcomes requires more than technologyit demands clean data, cross-functional collaboration, disciplined model lifecycle management, and sustained leadership commitment. But for organizations willing to make that investment, predictive maintenance transforms asset management from a necessary expense into a strategic capability that drives efficiency, resilience, and competitive advantage.

References

[1] OxMaint. "Predictive Maintenance vs Preventive Maintenance." OxMaint Blog. https://www.oxmaint.com/blog/post/predictive-maintenance-vs-preventive-maintenance

[2] Transatel. "Predictive vs Traditional Maintenance." Transatel FAQ. Machine Metrics. "Predictive vs Preventative Maintenance." MachineMetrics Blog. https://www.machinemetrics.com/blog/predictive-vs-preventative-maintenance

[3] IBM. "Predictive vs Preventive Maintenance." IBM Think. https://www.ibm.com/think/topics/predictive-vs-preventive-maintenance

[4] ServiceChannel. "Preventive vs Predictive Maintenance." ServiceChannel Blog. https://servicechannel.com/blog/preventive-vs-predictive-maintenance/

[5] AssetWatch. "Preventive vs Predictive Maintenance." AssetWatch Blog. https://www.assetwatch.com/blog/preventive-vs-predictive-maintenance

[6] Fogwing. "Preventive vs Predictive Maintenance." Fogwing Blog. https://www.fogwing.io/blog/preventive-vs- predictive/

[7] Opsima. "Preventive vs Predictive Maintenance." Opsima Blog. https://opsima.com/blog/operational-insights/preventive-vs-predictive-maintenance/

[8] Microsoft. "Sensor Data Intelligence Home Page (Preview)." Microsoft Learn, Dynamics 365 Documentation. https://learn.microsoft.com/en-us/dynamics365/supply-chain/sensor-data-intelligence/sdi-home-page

[9] CFB Solutions. "Dynamics 365: Asset Management." Datasheet (PDF). https://blog.cfbs-us.com/hubfs/Brochure/D365E/[Dynamics 365] SCM Datasheet - Asset management.pdf

[10] Ellipse Solutions. "Exploring Sensor Data Intelligence for Microsoft Dynamics 365 Supply Chain Management." Ellipse Solutions Blog. https://ellipsesolutions.com/factory-of-the-future-exploring-sensor-data- intelligence-for-microsoft-dynamics-365-supply-chain-management/

[11] Microsoft. "Asset Maintenance Scenario (Preview)." Microsoft Learn, Supply Chain Management. https://learn.microsoft.com/en-us/dynamics365/supply-chain/sensor-data-intelligence/sdi-scenario- asset-maintenance

[12] Key Dynamics Solutions. "Microsoft Dynamics 365 Supply Chain Management (D365 SCM)." Key Dynamics Solutions. https://keydynamicssolutions.com/d365-supply-chain-management/

[13] Microsoft. "Maintain Assets in Dynamics 365 Supply

Chain Management." YouTube. https://www.youtube.com/watch?v=GfnTrdzxpvE

[14] Dynamics365forOperations.de. "IoT Integrates with Microsoft Dynamics 365 Supply Chain." Blog post, November 2022. https://dynamics3654operations.de/2022/11/21/iot-goes-dynamics-365-supply-chain/

[15] Microsoft. "What Is Microsoft Fabric." Microsoft Learn, Fabric Fundamentals. https://learn.microsoft.com/en-us/fabric/fundamentals/microsoft-fabric-overview

[16] Microsoft. "End-to-End Data Platform, an Example Architecture." Azure Architecture Center. https://learn.microsoft.com/enus/azure/architecture/example-scenario/dataplate2e/data-platform-end- to-end

[17] Microsoft. "Microsoft FabricA Unified Analytics Platform for the AI Era." Microsoft Fabric. https://www.microsoft.com/en-us/microsoft-fabric

[18] Microsoft. "What Is Data Science in Microsoft Fabric?" Microsoft Learn, Fabric Data Science. https://learn.microsoft.com/en-us/fabric/data-science/data-science-overview

[19] Microsoft. "What's New in Microsoft Fabric." Microsoft Learn, Fabric Fundamentals. https://learn.microsoft.com/en-us/fabric/fundamentals/whats-new

[20] Microsoft Fabric Blog. "Microsoft Fabric Explained for Existing Synapse Users." Microsoft Fabric Blog. https://blog.fabric.microsoft.com/en-us/blog/microsoft-fabric-explained-for-existing-synapse-users

[21] Microsoft. "Small and Medium Business Modern Data Platform." Azure Solution Ideas. https://learn.microsoft.com/en-us/azure/architecture/solution-ideas/articles/small-medium-modern- data-platform

[22] Microsoft Fabric Blog. "From Azure Synapse and Azure Data Factory to Microsoft Fabric: The Next-Gen Analytics Leap." Microsoft Fabric Blog. https://blog.fabric.microsoft.com/en-US/blog/from-azure-synapse- and-azure-data-factory-to-microsoft-fabric-the-next-gen-analytics-leap

[23] Microsoft. "Predictive Maintenance for IoT." Azure Solution Ideas. https://learn.microsoft.com/en-us/azure/documentdb/solutions-iot

[24] Microsoft. "IoT and Data Analytics." Azure Solution Ideas. https://learn.microsoft.com/en-us/azure/architecture/solution-ideas/articles/iot-azure-data-explorer

[25] DataLabs. "Leveraging Azure AI for Predictive Maintenance in Industrial IoT." DataLabs. https://datalabs.io/leveraging-azure-ai-for-predictive-maintenance-in-industrial-iot/

[26] DT Engineering. "Predictive Maintenance with IoT Sensors." DT Engineering Articles. https://www.dtengineering.com/articles/predictive-maintenance-with-iot-sensors

[27] OneUptime Blog. "How to Ingest Real-Time IoT Hub Telemetry into Azure Digital Twins Using Azure Functions." February 2026. https://oneuptime.com/blog/post/2026-02-16-how-to-ingest-real-time-iot- hub-telemetry-into-azure-digital-twins-using-azure-functions/view

[28] Microsoft. "Predictive Maintenance Architecture." Microsoft Learn, Real-Time Intelligence. https://learn.microsoft.com/en-us/fabric/real-time-intelligence/architectures/predictive-maintenance

[29] Quadrant Resources. "Predictive Maintenance for Manufacturing." Microsoft AppSource. https://marketplace.microsoft.com/en-us/product/quadrantresourcellc.predictive_maintenance_manufacturing

[30] Dynatech Consultancy. "Powering Dynamics 365 Data Insights with Microsoft Fabric." Dynatech Blog. https://dynatechconsultancy.com/blog/powering-dynamics-365-data-insights-with-microsoft-fabric

[31] Acuvate. "Predictive Maintenance for Turbines with Microsoft Fabric & Azure ML." Acuvate Blog. https://acuvate.com/blog/predictive-maintenance-turbines-microsoft-fabric-azure-ml/

[32] Promethium. "What Is Microsoft Fabric? The Ultimate Guide." Promethium Guides. https://promethium.ai/guides/what-is-microsoft-fabric-guide/

[33] Techment. "Microsoft Fabric Architecture Explained. Techment Blog. https://www.techment.com/blogs/microsoft-fabric-architecture-explained/

[34] Digmatix. "Microsoft Fabric: A Complete Guide to the Unified Data Platform." Digmatix Blog. https://www.digmatix.com/en/blogs/microsoft-fabric-complete-guide-unified-data-platform

[35] Microsoft. "Deep Dive: Sensor Data Intelligence Add-in for Dynamics 365." YouTube. https://www.youtube.com/watch?v=OBIeXxkmJKI

[36] YouTube. "Predictive Maintenance Using MachineLearning and IoT." https://www.youtube.com/watch?v=xa0pvpR8OEg