

Predictive Maintenance for Server Failure in Virtual Environments

By

Wong Pei Kei

A REPORT

SUBMITTED TO

Universiti Tunku Abdul Rahman

in partial fulfillment of the requirements

for the degree of

BACHELOR OF INFORMATION TECHNOLOGY (HONOURS)

COMMUNICATIONS AND NETWORKING

Faculty of Information and Communication Technology

(Kampar Campus)

JUNE 2025

ACKNOWLEDGEMENTS

I would like to take this opportunity to thank all the people who have supported me and are still supporting me throughout the journey of researching and developing this project on predictive maintenance for server failure. I am truly grateful to my supervisor, Dr. Aun Yichiet, for being highly supportive of the project's direction and for providing valuable feedback that helped improve the effectiveness of the solution in addressing real-world problems. A special thanks goes to my parents, who have been incredibly supportive of my studies and have always made sure that I am in good health, both physically and mentally. Through this journey, I have come to realize how important our well-being is, as being in poor shape can seriously affect the quality of both the research and the final report. One key factor behind the success of this project is that the topic was personally chosen by me. Because of that, I felt a strong sense of responsibility and passion, dedicating countless days and nights to ensure the outcome met my expectations and standards. I genuinely hope that this report will be helpful to anyone working on a similar topic or looking for guidance in this area. Thank you very much, and I wish you the best of luck with your own research journey.

COPYRIGHT STATEMENT

© 2025 Wong Pei Kei. All rights reserved.

This Final Year Project proposal is submitted in partial fulfillment of the requirements for the degree of Bachelor of Information Technology (Honours) Communications and Networking at Universiti Tunku Abdul Rahman (UTAR). This Final Year Project proposal represents the work of the author, except where due acknowledgment has been made in the text. No part of this Final Year Project proposal may be reproduced, stored, or transmitted in any form or by any means, whether electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the author or UTAR, in accordance with UTAR's Intellectual Property Policy.

ABSTRACT

This project develops a predictive maintenance framework for server failures in virtual environments through a six-stage workflow. Large-scale datasets from Google Cluster, Backblaze HDD, CINECA M100, and Azure VM traces were modeled to establish domain-specific baselines, followed by the design of a unified schema enabling cross-domain integration and transfer learning. Evaluation confirmed CPU stress and thermal load as dominant predictors, with thresholding guided by operational risk. A simulation tool with calibrated mathematical risk models, probabilistic scoring, and interactive dashboards was implemented to address deterministic prediction issues, and a pseudo real-time system was demonstrated using streamed logs. The workflow delivers a complete, simulation-ready predictive maintenance pipeline with both academic rigor and practical deployment value.

Area of Study: Predictive Maintenance, Machine Learning

Keywords: Server Failure Prediction, Hardware Metrics, Real-time Telemetry, Data Analysis, Machine Learning

TABLE OF CONTENTS

ACKNOWLEDGEMENTS.....	II
COPYRIGHT STATEMENT	III
ABSTRACT	IV
TABLE OF CONTENTS.....	V
LIST OF FIGURES	IX
LIST OF TABLES	XI
LIST OF SYMBOLS	XII
LIST OF ABBREVIATIONS	XIV
CHAPTER 1	16
Introduction	16
1.1 Problem Statement and Motivation.....	18
1.2 Research Objectives	19
1.3 Project Scope and Direction.....	20
1.4 Contributions.....	22
1.5 Background Information	23
1.6 Report Organization	25
CHAPTER 2	26
Literature Reviews	26
2.1 Prior Arts in Predictive Maintenance for Server Failure	26

2.1.1 Predictive Maintenance in Server Infrastructure.....	26
2.1.2 Hardware Health Monitoring	27
2.1.3 Techniques for Failure Prediction	27
2.2 Datasets and Data Collection for Predictive Maintenance.....	29
2.2.1 Processor & Motherboard Datasets.....	29
2.2.2 Sensor & Hardware Metrics Datasets	29
2.2.3 Failure Type Datasets.....	30
2.2.4 Usage Pattern Datasets	30
2.3 Fact Finding and Data Analysis Techniques.....	31
2.3.1 Scientific Methods in PM Data Analysis	31
2.3.2 Data Element Analysis	32
2.3.3 Report Use Analysis.....	33
2.4 Existing Predictive Maintenance Models and Techniques	35
2.4.1 Machine Learning Models in PM	35
2.4.2 Deep Learning Models	35
2.4.3 Hybrid Models	36
2.5 Critical Review of Previous Works	37
2.5.1 Strengths of Existing Approaches	37
2.5.2 Weaknesses and Gaps	37
2.5.3 Comparison with Proposed Solution.....	38
CHAPTER 3	41
Proposed Method/Approach	41
3.1 System Development Models	41
3.1.1 System Development Model 1 – Waterfall Approach	41
3.1.2 System Development Model 2 – Spiral Model.....	41
3.1.3 System Development Model 3 – Agile Development.....	42
3.1.4 Selected Model.....	42
3.2 System Requirement	43
3.2.1 Hardware	43
3.2.2 Software	43
3.2.3 Data Processing Environment.....	45
3.3 Functional Requirements	45
3.4 Timeline	46
3.5 Estimated Cost	47
3.6 Implementation of Issues and Challenges.....	47
CHAPTER 4	50

System Design.....	50
4.1 System Architecture Diagram	50
4.2 Use Case Diagram.....	56
4.3 Activity Diagram.....	58
4.4 System Block Diagram	64
4.4.1 User	65
4.4.2 Virtualized Environment (Cloud, HPC, and Storage Systems)	65
4.4.3 Risk Modeling Engine.....	66
4.4.4 Simulation and Dashboard Service	66
4.4.5 Data Services.....	67
4.5 System Components.....	67
4.6 Dataset Summary Table	68
4.7 Sensor Metric to Hardware Component Mapping Diagram	70
4.8 Data Flowchart	71
4.9 System Equation	72
4.10 Functional Modules in the System.....	73
4.11 Data & Model Specification Summary	74
4.12 GUI Design	75
CHAPTER 5	76
System Implementation.....	76
5.1 Hardware Setup	76
5.2 Software Setup	77
5.3 Setting and Configuration	78
5.4 System Operation	81
CHAPTER 6	86
System Evaluation and Discussion	86
6.1 System Testing and Performance Metrics	86
6.2 Testing Setup and Result.....	90

6.3 Project Challenges.....	91
6.4 SWOT Analysis	92
6.5 Objectives Evaluation	93
6.6 Feasibility Analysis.....	94
6.7 Scalability Consideration	94
6.8 Future Work	95
CHAPTER 7	96
Conclusion	96
Recommendation.....	97
REFERENCES.....	98
APPENDIX A	106
A.1 Code Sample	106
A.2 Poster.....	107

LIST OF FIGURES

Figure Number	Title	Page
Figure 2.3.1	Sensor Data Flow	31
Figure 2.3.2	SQL Data Flow	33
Figure 3.2.1	Specifications of Laptop.	43
Figure 3.2.2	Google Colab.	43
Figure 3.2.3	Kaggle.	44
Figure 3.2.4	Streamlit.	44
Figure 3.2.5	Jupyter Notebook.	45
Figure 3.5	Timeline.	46
Figure 4.1.1	Data Unification and Model Training Pipeline.	51
Figure 4.1.2	Simulation Tool Architecture.	53
Figure 4.1.3	Pseudo Real-Time Monitoring Simulation.	54
Figure 4.2.1	Data Scientist Use Cases.	56
Figure 4.2.2	ML Engineer Use Cases.	56
Figure 4.2.3	System Admin Use Cases.	57
Figure 4.2.4	End User Use Cases.	57
Figure 4.3.1	Per-Dataset Baselines.	58
Figure 4.3.2	Unified Schema Development.	59
Figure 4.3.3	Model Training.	60
Figure 4.3.4	Model Evaluation.	61
Figure 4.3.5	Simulation Tool Development.	62
Figure 4.3.6	Pseudo Real-Time Monitoring.	63
Figure 4.4.1	Cloud Failure Prediction Service	64
Figure 4.4.2	HPC Thermal-Power Risk Service	65
Figure 4.4.3	Storage Health Monitoring Service	65
Figure 4.8	Predictive Maintenance Workflow	71
Figure 4.12	GUI Design	74

Figure 5.4.1	Terminal Execution of Streamlit app	81
Figure 5.4.2	Streamlit Dashboard	81
Figure 5.4.2	Streamlit Dashboard – domain selection panel	81
Figure 5.4.3	Incoming Batch Passes Through Models	82
Figure 5.4.4	Composite risk scores	82
Figure 5.4.5	Real-time risk gauge visualization	83
Figure 5.4.6	Historical Trend Analysis Chart	84
Figure 5.4.7	System provides actionable guidance	84
Figure 5.4.8	Prediction Records	85
Figure 6.1	System Performance - AUC & F1 Score	87
Figure 6.2	Confusion Matrices	87
Figure 6.3	ROC Analysis	88
Figure 6.4	Precision Recall	88
Figure 6.5	Feature Importance	89
Figure 6.6	Cross-Domain Performance	89
Figure 6.7	Temporal Validation	90

LIST OF TABLES

Table Number	Title	Page
Table 1.5	Key Terms and Definitions Used in Predictive Maintenance for Server Failure	24
Table 3.1	Selected Model Workflow	42
Table 3.2	Specifications of laptop	43
Table 3.3	Functional Requirements	45
Table 3.5	Estimated Cost	47
Table 4.2	Actor Interactions	57
Table 4.5	Algorithm of Core Pipeline	67
Table 4.6	Dataset Summary	68
Table 4.7	Mapping of Sensor Metrics to Hardware	70
Table 4.9	System Equation	72
Table 4.10	Functional Modules	73
Table 4.11	Data & Model Specification	74
Table 5.2	Software Requirements for System Implementation	77
Table 5.3	Model Configuration	79
Table 6.4	SWOT Analysis	92

LIST OF SYMBOLS

β	Constant coefficient (for CPU failure model)
γ	Constant coefficient (for Memory failure model)
T	Temperature (°C)
U_{cpu}	CPU utilization (%)
U_{ram}	RAM utilization (%)
t_{up}	System uptime (hours)
τ_{temp}	Temperature threshold (75°C)
τ_{cpu}	CPU utilization threshold (90%)
τ_{up}	Uptime threshold (5000 hours)
τ_{ram}	RAM usage threshold (85%)
V	Voltage (V)
V_{min}	Minimum safe voltage (11.5V)
V_{max}	Maximum safe voltage (12.5V)
$I(\cdot)$	Indicator function (returns 1 if condition is true, 0 otherwise)
P_{cpu}	Probability of CPU failure
P_{mem}	Probability of memory failure
P	Power (W)
τ	Torque (Nm)
ω	Angular velocity (rad/s)
N	Rotational speed (rpm)
R_{TV}	Temperature-voltage ratio (°C/V)
μ_t	Moving average at time t
σ_t	Standard deviation at time t
x_i	Observed value at time i
t	Current time index
$Risk(x)$	Risk classification level
$P(x)$	Model-predicted failure probability for observation x
θ	Model parameters

L	Cross-entropy loss function
$\Omega(\theta)$	Regularization term
f_{θ}	Model function parameterized by θ
y_i	True label of observation i
τ^*	Decision threshold
\hat{y}	Predicted failure label (1 or 0)
W_{τ}	Tool wear–torque interaction term
w	Tool wear (minutes)

LIST OF ABBREVIATIONS

<i>PM</i>	Predictive Maintenance
<i>OHM</i>	Open Hardware Monitor
<i>HWINFO</i>	Hardware Information Tool
<i>SMOTE</i>	Synthetic Minority Over-sampling Technique
<i>ML</i>	Machine Learning
<i>AI</i>	Artificial Intelligence
<i>ANN</i>	Artificial Neural Network
<i>LSTM</i>	Long Short-Term Memory
<i>AUC</i>	Area Under Curve
<i>CPU</i>	Central Processing Unit
<i>GPU</i>	Graphics Processing Unit
<i>RAM</i>	Random Access Memory
<i>NIC</i>	Network Interface Card
<i>IC</i>	Integrated Circuit
<i>S.M.A.R.T.</i>	Self-Monitoring, Analysis, and Reporting Technology
<i>CSV</i>	Comma Separated Values
<i>RUL</i>	Remaining Useful Life
<i>SVM</i>	Support Vector Machine
<i>API</i>	Application Programming Interface
<i>ROC</i>	Receiver Operating Characteristic

CHECKLIST

TICK (√)	DOCUMENT ITEMS
	Your report must include all the items below. Put a tick on the left column after you have checked your report with respect to the corresponding item.
√	Title Page
√	Copyright Statement
√	Acknowledgement
√	Abstract (with Area of Study and Keywords)
√	Table of Contents
√	List of Figures (if applicable)
√	List of Tables (if applicable)
√	List of Symbols (if applicable)
√	List of Abbreviations (if applicable)
√	Chapters / Content
√	References
√	All references are cited in the report, especially in the chapter of literature review
√	Appendices (if applicable)
√	Poster
√	I agree 5 marks will be deducted due to incorrect format, declare wrongly the ticked of these items, and/or any dispute happening for these items in this report.

CHAPTER 1

Introduction

As businesses increasingly depend on virtualized infrastructures for cloud computing, data handling, and day-to-day operations, maintaining server reliability has become more important than ever. These infrastructures use technologies like virtual machines and resource pooling to boost efficiency and scalability [1], [2]. However, as systems grow more complex—with different types of processors, workloads, and environments—keeping them running smoothly becomes harder, especially across diverse setups [7].

Unexpected server failures can cause serious issues, especially in industries like finance, telecommunications, and cloud services, where uptime directly impacts income and customer satisfaction [3], [4]. For instance, a single server outage could cost a company around \$5,600 for every minute it's down, showing how essential it is to have a maintenance plan that can stop problems before they happen [3].

One effective solution is predictive maintenance (PM), which uses artificial intelligence (AI) and machine learning (ML) to spot signs of potential failures early on by analyzing real-time and historical data [5], [6]. These systems pull information from sensors, system logs, and workload data to predict when something might go wrong and help prevent it with timely actions.

In FYP1, the project introduced an AI-based predictive maintenance system by exploring multiple datasets such as Laptop Motherboard Health Monitoring, Microsoft Azure Predictive Maintenance, C-MAPSS, AI4I 2020, Machine Failure using Sensor Data, Hard Drive Test Data, and Google Cluster Workload [8]. These datasets provided a starting point for building models that learned from various processor types, hardware indicators, and usage patterns. To address the lack of real-world server failure records, simulated and short-term telemetry data were used for validation.

Building on this foundation, FYP2 advances the research with a structured six-stage workflow that shifts the focus from broad exploratory datasets to domain-specific and unified predictive modeling. This includes integrating domain-specific datasets (Google Cluster, Backblaze, CINECA, Azure VM) into a unified schema, training

CHAPTER 1

cross-domain models, and developing a Streamlit-based simulation tool for pseudo real-time monitoring.

Unlike FYP1, which relied on synthetic augmentation and real-time telemetry collection tools (HwiNFO, OHM, Zabbix Agent), FYP2 emphasizes deployability and cross-domain generalization, making the predictive maintenance pipeline lightweight, reproducible, and adaptable across virtualized environments. The final system is not only academically rigorous but also practically deployable, offering a simulation-ready dashboard for early fault detection and proactive server management.

1.1 Problem Statement and Motivation

Problem Statement

As modern organizations expand their reliance on virtualized and cloud-based infrastructures, ensuring continuous server uptime has become increasingly complex. These environments often involve a wide range of server configurations, including different processor architectures (Intel, AMD, Apple, etc.), varying workloads, and diverse hardware setups. Traditional maintenance strategies, such as periodic checks or reactive repairs, are no longer sufficient in such dynamic environments. They often fail to detect early warning signs of system failure, leading to unexpected downtime that can severely impact business continuity and service availability.

Existing predictive maintenance (PM) solutions remain limited. Many are tailored to specific hardware, lack real-time adaptability, or rely on narrowly focused datasets that do not generalize across systems. This results in inaccurate predictions or failure to detect issues in heterogeneous environments. Furthermore, many organizations lack access to an integrated PM system that combines diverse datasets, real-time data, and scalable AI models. This gap highlights the need for a flexible, adaptable, and accurate predictive maintenance solution for server infrastructure management.

Motivation

The motivation behind this project lies in the growing demand for smarter predictive maintenance systems that can minimize downtime, optimize resource usage, and improve reliability across diverse computing environments. As cloud and edge infrastructures scale, their complexity requires solutions beyond static monitoring. A robust AI-driven PM model that leverages both historical datasets and simulated real-time data offers a proactive approach to addressing this need. From an academic and personal perspective, this project provides an opportunity to apply AI and ML practically in system maintenance—an area often overlooked in IT operations. By building a reproducible pipeline and a deployable dashboard, this work aims to reduce unplanned outages and contribute to the development of intelligent infrastructure management.

1.2 Research Objectives

The main goal of this project is to design and develop an AI-driven predictive maintenance system that can accurately predict server failures across heterogeneous environments. Unlike FYP1, which mainly explored diverse datasets and real-time telemetry collection, FYP2 emphasizes structured workflows, domain-specific integration, and deployable simulation tools. The system aims to proactively reduce downtime, optimize resource utilization, and support scalable server monitoring.

The project will focus on integrating and processing data from multiple domain-specific datasets such as Google Cluster, Azure VM, Backblaze, and CINECA HPC. These datasets are standardized into a unified schema, allowing the development of both domain-specific and cross-domain models. Instead of relying on physical real-time sensors, pseudo real-time monitoring will be achieved through a Streamlit-based dashboard that simulates live server conditions. This approach makes the system more lightweight, reproducible, and adaptable across virtualized environments.

This project will specifically aim to:

1. **Develop a unified machine learning pipeline** in Python that supports data preprocessing, schema integration, feature engineering, and model training across multiple domains.
2. **Train and validate predictive models** capable of achieving at least 80% accuracy, evaluated through precision, recall, F1-score, and AUC metrics.
3. **Simulate real-time monitoring** by replaying historical traces in a Streamlit dashboard with risk classification outputs and interactive visualization.
4. **Ensure adaptability and deployability** by keeping the system lightweight and reproducible, suitable for integration into different server environments.

However, the project will not cover areas such as direct log analysis, deep learning-based architectures, or large-scale enterprise deployment. Continuous real-time data ingestion and advanced scalability features are considered potential future extensions rather than immediate deliverables.

In short, this project aims to bridge the gap between multi-domain server data and intelligent decision-making using AI, delivering a deployable prototype that is both academically rigorous and practically useful.

1.3 Project Scope and Direction

Project Scope

This project aims to develop an AI-based predictive maintenance system capable of predicting server failures before they occur. The system leverages both historical datasets and pseudo real-time telemetry to provide early warnings and actionable insights for system administrators. To keep the project feasible within the final year project timeline, several scope limitations and clear deliverables are defined:

- **Target Hardware:** Focus on common x86 servers (Intel and AMD). Other processor types, such as Apple ARM or Nvidia, are included only in simulated datasets to test model flexibility.
- **Data Sources:** Uses public datasets including Google Cluster, Azure VM, Backblaze, and CINECA HPC. Historical traces are replayed to simulate real-time monitoring in the Streamlit dashboard.
- **Real-Time Simulation:** The system simulates live monitoring by feeding pre-recorded telemetry in timed intervals. Tools like HWiNFO, Open Hardware Monitor (OHM), and Zabbix Agent provide the reference for sensor metrics (temperature, voltage, fan speed) exported in CSV format.
- **Machine Learning Pipeline:** Implemented in Python, using Kaggle or Jupyter Notebook, covering preprocessing, feature engineering, model training, evaluation, and visualization.
- **Key Features:**
 1. Predict server failures with at least 80% accuracy using standard metrics (precision, recall, F1-score, AUC).
 2. Simulate pseudo real-time monitoring with interactive dashboards and auto-refreshing risk visualization.
 3. Provide outputs in human-readable formats (HTML, JSON) for proactive maintenance.
 4. Train and validate models on multiple domain datasets to ensure reliability and versatility.

CHAPTER 1

The project does not focus on building an enterprise-scale system, processing system logs, or integrating advanced visual diagnostics. These are considered future improvements.

Direction

This project serves as the foundation for a more advanced predictive maintenance system suitable for real-world deployment in data centers or cloud platforms. While the current version simulates real-time monitoring in a controlled environment, future enhancements could include:

- Continuous monitoring with live telemetry streams from servers or cloud services.
- Integration of additional input data such as logs, network metrics, or thermal imaging.
- Application of advanced AI methods, including deep learning and ensemble models, to improve prediction accuracy.
- Deployment using containers (Docker) or cloud-based frameworks for scalability.
- Enhanced user interfaces and potential multilingual support for broader applicability.

Overall, the project direction aims to develop a smart, adaptable system that minimizes server downtime, enhances reliability, and lays the groundwork for scalable predictive maintenance across heterogeneous IT environments.

1.4 Contributions

The contributions of this project aim to deliver an effective, intelligent, and practical solution for predicting server failures across diverse computing environments. By implementing a machine learning-based predictive maintenance system with a structured six-stage workflow, this project assists IT teams in minimizing downtime, reducing operational disruptions, and enhancing system reliability. The main contributions are as follows:

1. **Accurate and Early Failure Prediction**

The system analyzes historical datasets (Google Cluster, Azure VM, Backblaze, CINECA HPC) integrated into a unified schema to detect anomalies and forecast potential failures. Models are trained with cross-domain data and balanced using techniques like SMOTE to handle rare events. The predictive pipeline achieves high accuracy, precision, and recall, enabling administrators to respond proactively and prevent costly server outages.

2. **Deployable Simulation and Visualization**

To demonstrate real-world applicability without full live deployment, the system includes a Streamlit-based pseudo real-time dashboard. Historical telemetry traces are replayed in sliding windows to simulate live monitoring, with interactive risk gauges, trend charts, and parameter sliders. Outputs are provided in intuitive formats such as HTML dashboards and JSON reports, making it practical for IT teams to monitor server health and integrate insights into existing management tools.

3. **Cross-Domain Generalizability and Scalability**

By combining domain-specific datasets into a unified modeling framework, the system generalizes across Cloud, HPC, and Storage environments. While real-time simulation currently focuses on x86-based servers, the modular architecture allows future expansion to other processor types and hardware platforms. This design ensures adaptability and scalability for diverse infrastructures, demonstrating potential for broader deployment in data centers or enterprise networks.

1.5 Background Information

In today's digital era, server infrastructure forms the backbone of nearly all modern enterprises, cloud services, and data-intensive applications. Ensuring the reliability and uptime of servers is critical, as unexpected hardware failures can result in costly downtime, data loss, and service interruptions. Traditional maintenance strategies, such as reactive maintenance (fixing components only after failure) and scheduled preventive maintenance, are often inefficient or insufficient in complex, multi-environment systems. This has led to the rise of Predictive Maintenance (PM) — a data-driven approach that anticipates and prevents failures before they occur. PM leverages machine learning models to analyze historical and pseudo real-time telemetry data, enabling early detection of hardware degradation or failure risks.

Predictive maintenance techniques originally gained traction in manufacturing, industrial machinery, and aerospace (e.g., the C-MAPSS turbofan engine dataset). Applying these methods to IT infrastructure and server hardware is relatively new. Server environments present unique challenges, such as diverse processor types, complex telemetry (temperature, voltage, power usage, fan speed, workload patterns), and multiple failure types across hardware, software, and network components. Recent advancements in telemetry monitoring tools (HWiNFO, Open Hardware Monitor, Zabbix Agent) enable detailed and continuous server health data collection. However, real-world server failure datasets remain scarce and sensitive, limiting research progress.

FYP2 builds on FYP1 by implementing a structured six-stage predictive maintenance workflow that integrates domain-specific datasets (Google Cluster, Azure VM, Backblaze, CINECA HPC) into a unified schema. The project emphasizes cross-domain modeling, pseudo real-time simulation using Streamlit dashboards, and deployable machine learning pipelines. Key technical concepts include feature engineering (e.g., rolling statistics, workload intensity ratios, thermal and power metrics), class balancing (SMOTE), domain-specific and unified model training, risk classification, and simulation-ready dashboard visualization. The system shifts maintenance strategies from reactive or purely scheduled to proactive, enabling early interventions, reduced downtime, and improved server reliability.

This report guides the reader through both the theoretical foundation and practical implementation of the system, explaining how predictive maintenance principles are

CHAPTER 1

applied to heterogeneous server environments, even for readers without deep experience in server monitoring or machine learning.

Term	Definition
Predictive Maintenance (PM)	A maintenance strategy that uses data analysis tools and techniques to detect anomalies and predict equipment failures before they occur.
Telemetry Data	Automatic collection and transmission of measurements and operational data from hardware devices to a central system for monitoring and analysis.
Open Hardware Monitor (OHM)	A free and open-source tool that monitors temperature sensors, fan speeds, voltages, and load of hardware components in a computer system.
HWINFO	A comprehensive hardware monitoring and diagnostic tool that provides in-depth telemetry on CPU, GPU, memory, drives, and motherboard sensors.
Zabbix Agent	A lightweight software component that collects monitoring data (e.g., CPU usage, disk space, network activity) from servers and sends it to a Zabbix monitoring server.
Failure Types	Categories of possible system failures, such as hardware faults (CPU, disk), software crashes, and network disruptions.
Feature Engineering	The process of transforming raw data into meaningful input features that improve the performance of machine learning models.
SMOTE (Synthetic Minority Over-sampling Technique)	A method used to balance class distribution in imbalanced datasets by generating synthetic samples of the minority class.
Concept Drift	A phenomenon where the statistical properties of target variables change over time, potentially degrading model performance if not updated.
Hyperparameter Tuning	The process of optimizing parameters that govern the learning process of machine learning models to achieve the best performance.

Table 1.5 Key Terms and Definitions Used in Predictive Maintenance for Server Failure

1.6 Report Organization

The details of this research are presented in the following chapters. Chapter 2 reviews the background on predictive maintenance, relevant technologies, and existing systems, highlighting key hardware, software, and algorithm considerations. Chapter 3 outlines the methodology for developing the predictive maintenance system, including system development models, functional and technical requirements, project milestones, and cost estimation. Chapter 4 discusses the system design, covering architecture, functional modules, system flow, and GUI considerations. Chapter 5 presents the system implementation, detailing hardware and software setup, configuration, and operational workflow with supporting screenshots. Chapter 6 evaluates the system through testing, performance metrics, challenges, and objectives assessment. Finally, Chapter 7 concludes the project, summarizing contributions, limitations, and providing recommendations for future work.

CHAPTER 2

Literature Reviews

2.1 Prior Arts in Predictive Maintenance for Server Failure

2.1.1 Predictive Maintenance in Server Infrastructure

Predictive maintenance (PM) has become a pivotal approach for ensuring the reliability and efficiency of server infrastructure, particularly in large-scale data centers and cloud platforms. By leveraging real-time monitoring and intelligent data analysis, PM shifts maintenance practices from reactive or scheduled to proactive, significantly reducing downtime and costs.

Many leading organizations, including Google, AWS, and Microsoft Azure, have integrated AI-driven predictive models into their infrastructures to detect anomalies early and enhance fault resilience [13]. These implementations utilize machine learning (ML) and deep learning models—such as LSTM, Transformers, and Artificial Neural Networks (ANN)—to forecast failures in servers, storage systems, and network components using telemetry data like CPU usage, memory load, and disk I/O [13], [14], [15]. ANN, in particular, has demonstrated high accuracy in predicting hardware failures and estimating remaining useful life (RUL) [14].

Studies also highlight the application of evolving fuzzy-rule-based systems and hybrid ML techniques, such as One-Class SVMs and Elastic Stack Suites, for analyzing log streams and clustering anomalies [10], [11]. These methods enhance early fault detection and support preventive maintenance strategies. Moreover, the rise of edge computing allows part of the predictive workload to be processed near the data source, reducing latency and enabling faster responses to impending issues [12].

In summary, modern predictive maintenance frameworks combine diverse AI and ML techniques with edge computing and telemetry systems to deliver intelligent, self-healing infrastructures. These innovations are transforming server maintenance into a strategic function crucial for the resilience and scalability of mission-critical environments.

2.1.2 Hardware Health Monitoring

Monitoring hardware health is central to predictive maintenance strategies in server environments, where failures in components like CPUs, memory, disks, or GPUs can result in substantial operational costs. Academic studies underscore the need for real-time surveillance of system parameters to detect early signs of degradation. For instance, Decker et al. [10], [11] and Kumar [12] emphasize AI-assisted monitoring using server logs and sensor data, which improves fault detection by identifying anomalies in performance metrics such as CPU usage, memory load, and disk activity.

Telemetry systems are instrumental in gathering the data necessary for training predictive models. Research by Polu [13], BOUROGA and SAHRAOUI [14], and Gao et al. [15] collectively highlights how telemetry from critical resources informs machine learning algorithms for failure forecasting. These systems contribute to enhanced accuracy and responsiveness in detecting potential breakdowns.

On the practical side, Chapman [16] and Aldea et al. [17] provide valuable perspectives on commonly used monitoring tools like HWiNFO, Open Hardware Monitor (OHM), and Zabbix. These utilities collect detailed hardware metrics and are often integrated with higher-level platforms to support proactive decision-making.

Edge computing also plays a growing role in accelerating hardware health analytics. Mourtzis et al. [18] describe an edge-computing platform leveraging 5G and machine learning to minimize latency and improve fault prediction, further supporting real-time decision-making and reducing downtime.

In essence, effective hardware health monitoring combines academic methodologies, real-time telemetry, and industry tools within intelligent architectures that underpin predictive maintenance and ensure sustained server performance.

2.1.3 Techniques for Failure Prediction

Failure prediction in server systems has undergone a transformation from traditional statistical methods to sophisticated AI-driven models. Early techniques, rooted in reliability engineering and statistical modeling—such as multivariate regression,

ARIMA, and Hidden Markov Models (HMM)—focused on predefined thresholds and historical failure trends [13], [14]. While foundational, these approaches often struggle in dynamic and heterogeneous data center environments.

Machine learning (ML) techniques introduced a more adaptive framework. Models like Decision Trees, Support Vector Machines (SVM), and Random Forests analyze telemetry and log data to identify early warning signs of system degradation. Both supervised and unsupervised approaches enable these models to handle varied operational scenarios, improving fault detection accuracy and response [11], [13], [14]. Clustering methods further support anomaly detection when labeled data is sparse.

The progression to deep learning (DL) has further enhanced predictive capabilities. Convolutional Neural Networks (CNNs), Long Short-Term Memory (LSTM) networks, and Transformer-based models can extract complex temporal and spatial relationships from high-dimensional telemetry datasets [12], [13], [15]. These architectures are particularly effective in detecting subtle, evolving patterns that precede failures, making them well-suited for real-time and large-scale environments.

Hybrid and fuzzy-logic-based systems represent another evolution, offering interpretability and adaptability. By combining evolving rule-based mechanisms with learning algorithms, these approaches manage uncertainty and update models continuously based on live data. Their use in real-time anomaly detection demonstrates practical effectiveness in mission-critical settings [10], [12], [18].

Collectively, the literature reveals a clear trend: from rule-based systems toward intelligent, self-adaptive models capable of anticipating failures in complex infrastructures. This evolution supports more resilient and autonomous server maintenance strategies, ultimately driving efforts toward achieving zero-downtime operations.

2.2 Datasets and Data Collection for Predictive Maintenance

2.2.1 Processor & Motherboard Datasets

The Laptop Motherboard Health Monitoring Dataset by Zia [20] offers a synthetic dataset focused on detecting hardware issues in laptop motherboards. It includes key features like CPU usage, voltage, temperature, RAM, and fan speed across 2000 records labeled with specific failure types. While useful for educational experiments and machine learning prototyping, the dataset lacks real-world noise and is not representative of enterprise server environments.

Another widely used resource is the Microsoft Azure Predictive Maintenance Dataset [21], which consists of telemetry, error logs, and maintenance history for 100 industrial machines. It has been employed to model fault prediction using multivariate time-series data. Azab et al. [19] utilized this dataset in their discrete-event simulation model to optimize scheduling and maintenance in industrial systems. However, the dataset is tailored to manufacturing rather than IT hardware, and its application to server-level predictive maintenance is indirect at best.

2.2.2 Sensor & Hardware Metrics Datasets

The C-MAPSS Dataset from NASA [22] simulates engine degradation over time using sensor data across different operational conditions. It is ideal for remaining useful life (RUL) prediction and widely adopted in benchmark studies. However, it does not represent server hardware metrics, limiting its relevance to IT predictive maintenance.

Similarly, the AI4I 2020 Predictive Maintenance Dataset [23] from UCI is a synthetic dataset for tool failure prediction using process temperature, torque, and wear data. Though well-labeled and useful for testing classification models, it is based on a manufacturing machine and offers no direct applicability to server components.

The Machine Failure Prediction Dataset by Naeem [24] introduces multiple sensor types, including VOC, ultrasonic, and environmental measurements. While versatile for general machine failure modeling, it lacks domain-specific details for server infrastructure.

2.2.3 Failure Type Datasets

The Backblaze Hard Drive Dataset [25] is a rare real-world dataset that provides daily S.M.A.R.T. metrics and failure records for thousands of hard drives in data center environments. It enables long-term modeling of drive reliability but is limited to storage components, with no data on CPUs, motherboards, or network devices. Currently, there is a significant lack of publicly available datasets on network hardware failures such as switches and NICs. Sensitive operational data and proprietary concerns likely hinder dataset release in this domain, marking it as a key research opportunity.

2.2.4 Usage Pattern Datasets

Usage pattern datasets capture how systems are operated over time, which can provide vital context for predicting failures. One notable dataset is the Google Cluster Trace [26], which contains over 29 days of usage data from Google's data centers, including job arrival times, resource usage, and machine events. It has been used in workload modeling and anomaly detection research. While rich in operational detail, it lacks specific hardware failure labels, making it more suitable as a supplementary dataset for correlating workload patterns with maintenance events.

2.3 Fact Finding and Data Analysis Techniques

2.3.1 Scientific Methods in PM Data Analysis

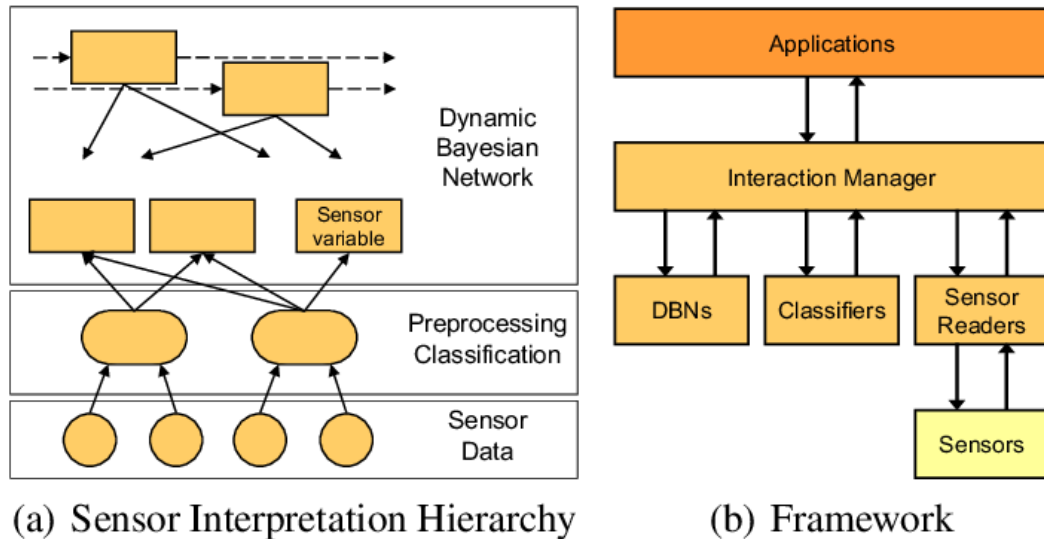


Figure 2.3.1: Sensor Data Flow

In predictive maintenance (PM), applying scientific and structured methods is critical to ensure accurate analysis of system health. Traditional practices often involve following well-defined maintenance procedures and system documentation, but modern techniques integrate advanced models such as Dynamic Bayesian Networks to interpret sensor data in real-time, as shown in Figure 2.3.1. Brandherm and Schmitz [27] proposed a modular framework that utilizes these networks for interpreting sensor data on mobile devices, highlighting how structured probabilistic models can support data-driven diagnostics. While the study focuses more on mobile systems, the methodology demonstrates how structured data interpretation methods can be effectively applied in PM scenarios. Additionally, Carvalho et al. [28] conducted a comprehensive review of machine learning (ML) techniques used in predictive maintenance, emphasizing the systematic process of data acquisition, preprocessing, and model development. Their study reinforces the importance of a scientific approach in data analysis to enable proactive fault detection and minimize equipment downtime. These techniques represent a modern evolution of traditional manual diagnostics, aligning well with the digital transformation goals in server infrastructure maintenance.

2.3.2 Data Element Analysis

Analyzing specific hardware components such as CPU, GPU, and RAM plays a central role in predictive maintenance, particularly in data center environments where performance monitoring is crucial. Several studies have highlighted the importance of evaluating these metrics to detect early signs of failure. For example, Decker et al. [10], [11] developed an evolving fuzzy-rule-based approach for real-time anomaly detection using log data, showing that tracking CPU and RAM performance can lead to effective predictive models. System architecture and data flow diagrams are often employed in these studies to visualize interactions between components, aiding in the interpretation of complex systems. Kumar [12] and Polu [13] further demonstrated that sensor data, when mapped through structured architectural diagrams, can identify bottlenecks and fault-prone components. Moreover, Bourouga and Sahraoui [14] applied multivariate regression techniques to telemetry data, emphasizing the predictive power of monitoring physical parameters that mirror server-side metrics like voltage and thermal output. Studies by Gao et al. [15] and Aldea et al. [17] also confirm that real-time monitoring of hardware resources combined with structured input-output analysis is essential for ensuring reliable performance. The consistent use of architecture diagrams and metric-based logging underlines a common framework for interpreting hardware health and forecasting failures in server ecosystems.

2.3.3 Report Use Analysis




Figure 2.3.2: SQL Data Flow

Diagnostic reports generated by tools such as HWPiNFO, Open Hardware Monitor (OHM), and Zabbix play a vital role in predictive maintenance (PM) by enabling continuous monitoring and analysis of system performance metrics. These tools provide detailed insights into hardware resource usage, environmental conditions, and potential system anomalies, supporting proactive failure prevention strategies. Zabbix, in particular, has seen enhancements in its recent versions, including machine learning integration and predictive functions that allow automated incident detection and forecasting based on monitored data trends [29]. Reports generated through Zabbix enable system administrators to identify early warning signs of hardware issues, optimize resource allocation, and plan maintenance activities before critical failures occur [30]. Figure 2.3.2 illustrates the typical data flow for such diagnostic tools, from raw HTTP metrics to processed insights.

Moreover, the use of automated reporting and root cause analysis within Zabbix helps streamline troubleshooting processes, improving the overall reliability and efficiency of IT infrastructure. Zabbix also allows for the integration of diverse data sources, making it effective for comprehensive system analysis. Such data can be used not only for operational decision-making but also to inform stakeholders about potential risks and the general health of the infrastructure [29], [30].

Complementing these capabilities, tools like HWPiNFO are crucial for environmental monitoring. They continuously log electrical parameters and thermal conditions, which

are critical factors influencing server hardware performance and longevity. Reports from these tools can reveal trends, such as rising component temperatures or voltage irregularities, that may indicate a developing issue [31]. Additionally, diagnostic stress testing using tools like AIDA64 and Prime95 enables engineers to assess component behavior under load, supporting the identification of weak points that are susceptible to failure. This kind of testing and analysis helps determine the reliability of CPUs, GPUs, and memory modules under extreme operational scenarios [31].

Ultimately, combining stress testing, environmental monitoring, and automated diagnostic reporting provides a multi-faceted approach to predictive maintenance. These tools help technicians detect abnormal patterns early, understand root causes of failures, and schedule preventive actions effectively. Such integration is essential for maintaining server uptime, avoiding costly downtimes, and ensuring long-term system resilience.

2.4 Existing Predictive Maintenance Models and Techniques

2.4.1 Machine Learning Models in PM

Machine learning (ML) techniques have played a significant role in advancing predictive maintenance (PM) strategies in server and data center environments. Commonly applied models such as Support Vector Machines (SVM), Random Forest (RF), and Gradient Boosting algorithms have shown effectiveness in identifying early signs of hardware or system failures by analyzing telemetry and log data. For example, studies by Decker et al. [11], Polu [13], and Gao et al. [15] demonstrated how supervised learning approaches like SVM and RF can accurately predict machine failures in data centers. These models excel at handling structured datasets and offer interpretable outcomes, making them suitable for applications where explainability is essential. Additionally, Mourtzis et al. [18] highlighted the integration of ML models with edge computing platforms to calculate Remaining Useful Life (RUL) of components, further enhancing real-time decision-making capabilities. However, these models can face limitations when dealing with highly dynamic environments or unstructured data, and their performance often depends heavily on the quality and completeness of the input features.

2.4.2 Deep Learning Models

Deep learning (DL) models have emerged as powerful alternatives to traditional ML methods, especially for analyzing complex, high-dimensional data such as time-series telemetry and system logs. Models like Long Short-Term Memory (LSTM) networks and Transformer-based architectures are particularly adept at capturing temporal dependencies and subtle failure patterns over time. Studies by Kumar [12], Polu [13], and Gao et al. [15] demonstrated how LSTM networks can effectively learn from sequential data to enhance fault prediction accuracy. Transformer-based models have also been explored to improve performance in predictive maintenance tasks [13], [15]. Despite their superior ability to model complex data relationships, a notable challenge with deep learning approaches is the issue of explainability. Unlike traditional ML models, deep neural networks often operate as "black boxes," making it difficult to

interpret the rationale behind their predictions—a drawback that can hinder their adoption in critical infrastructure contexts where transparency is vital.

2.4.3 Hybrid Models

Hybrid models, which combine multiple data sources and analytical techniques, have gained traction in predictive maintenance due to their ability to leverage the strengths of different methods. These models typically integrate sensor data, environmental data, and workload or log data to refine predictions and improve system diagnostics. Studies by Decker et al. [11], Polu [13], and Mourtzis et al. [18] illustrated how combining diverse data streams can enhance the accuracy and robustness of predictive maintenance systems. Furthermore, Kumar [12] emphasized the role of edge computing in hybrid frameworks, which enables real-time analytics by processing data closer to its source. A practical example of hybrid model application can be found in the work of Azab et al. [19], where a pharmaceutical company case study validated the integration of ML-assisted simulation with sensor and environmental data for dynamic scheduling and predictive maintenance. Hybrid approaches are increasingly being recognized for their capacity to balance predictive accuracy, computational efficiency, and contextual adaptability in real-world industrial and server environments.

2.5 Critical Review of Previous Works

2.5.1 Strengths of Existing Approaches

Existing predictive maintenance (PM) approaches have demonstrated several notable strengths, particularly in terms of real-world applicability and model robustness. Several studies have successfully implemented PM frameworks in actual data center and industrial environments, illustrating the feasibility of their deployment beyond theoretical research. For instance, Decker et al. [10] and Decker de Sousa et al. [11] showcased machine learning and fuzzy-rule-based systems capable of performing real-time anomaly detection in data centers, demonstrating their potential for practical application. Additionally, many models presented in the literature are built on robust datasets, which contribute to their high accuracy and reliability. Techniques such as deep learning models (e.g., LSTM and Transformers) and advanced regression methods have been effective in improving predictive performance [13], [15]. Furthermore, approaches combining edge computing with AI, as explored by Kumar [12] and Mourtzis et al. [18], have shown the promise of reducing latency and enhancing real-time data processing capabilities, making these solutions highly relevant for modern data-intensive environments. The integration of machine learning into scheduling strategies, such as in dynamic flow-shop scheduling [19], also highlights how PM can be embedded into broader operational workflows, further enhancing system efficiency and equipment uptime.

2.5.2 Weaknesses and Gaps

Despite the strengths, there are several notable weaknesses and gaps in the current PM literature. One of the most critical limitations is the lack of server-specific predictive maintenance models. Much of the existing research tends to focus on industrial equipment, turbofan engines, or general data center infrastructure, with limited emphasis on server hardware components such as processors and motherboards [10], [11]. Moreover, many datasets used in these studies lack diversity and are often constrained to specific environments or case studies. For example, some works are limited to pharmaceutical production lines or single data centers, restricting the generalizability of the models [19].

Another prominent gap lies in scalability and adaptability. Many approaches, including those based on fuzzy-rule systems [10] and deep learning models [13], face challenges when handling large-scale, dynamic cloud or data center environments. Ensuring real-time processing without performance bottlenecks remains difficult across these studies [12], [15]. Additionally, the heavy reliance on high-quality and comprehensive datasets is a recurring issue. The effectiveness of many PM models diminishes when data is noisy, incomplete, or varies significantly between different systems [14], [15].

Explainability is also a persistent concern. While advanced AI models can achieve high accuracy, their decision-making processes are often opaque, making it difficult for administrators to trust and interpret predictions [13]. Finally, the resource-intensive nature of most approaches—both in terms of computational requirements and the need for expert knowledge to implement and maintain systems—poses a significant barrier to adoption, particularly in smaller organizations with limited budgets and expertise [12], [18].

2.5.3 Comparison with Proposed Solution

While existing predictive maintenance (PM) frameworks have demonstrated substantial advancements, particularly in large-scale server infrastructure, the proposed solution in this project offers several key improvements that aim to address the identified gaps in the literature. Previous studies have effectively integrated AI-driven models, such as LSTM, Transformers, and Artificial Neural Networks (ANN), to forecast failures by analyzing telemetry data like CPU usage, memory load, and disk activity [13], [14], [15]. However, many of these works rely on relatively narrow datasets that are often limited to single environments or specific case studies, which reduces model generalizability [19]. Furthermore, much of the focus has been on data center-wide infrastructures or industrial machinery, with less emphasis on individual server components such as processors, motherboards, and sensor-driven metrics [10], [11].

In contrast, the proposed workflow strategically integrates diverse datasets that encompass various processor types, hardware metrics, failure types, and usage patterns. The inclusion of datasets spanning from the Laptop Motherboard Health Monitoring Dataset and Microsoft Azure Predictive Maintenance Dataset, to C-MAPSS Turbofan

Engine Simulation and Google 2019 Cluster sample, ensures that the model captures a wide variety of configurations, operational scenarios, and failure behaviors. This broad dataset diversity enhances the robustness and transferability of the predictive models compared to those relying on a single source or homogeneous environment.

Another novel aspect of the proposed solution lies in its feature engineering process. While existing studies have used conventional telemetry data like CPU load and memory usage [13], [14], [15], this project extends feature engineering by incorporating derived metrics, rolling statistics, and hybrid features (such as temperature-voltage ratios and CPU-RAM interactions). These complex features aim to uncover deeper temporal and cross-component relationships, which can improve the early detection of hardware degradation. Previous works have acknowledged the benefit of evolving fuzzy-rule-based systems and hybrid ML models for anomaly detection [10], [11], but they rarely emphasize advanced, interpretable feature synthesis at this granularity.

Moreover, although current industry-standard monitoring tools like HWiNFO, Open Hardware Monitor, and Zabbix have been effectively integrated in practical settings for hardware health monitoring [16], [17], the proposed solution goes a step further by combining these real-time monitoring utilities with an end-to-end predictive model pipeline. This integration allows the system not only to monitor but also to analyze historical and near-real-time data, simulate different failure types, and provide actionable risk assessments. Existing literature also highlights the increasing role of edge computing for reducing latency in hardware health analytics [12], [18]; similarly, the proposed framework leverages real-time data streams and lightweight models to enable faster failure predictions directly from live data.

Finally, the comprehensive deployment readiness of the proposed system — with modular pipelines, risk classification outputs, interpretability modules, and example inference scripts — enhances its practicality and usability in real-world settings. This level of systematization and documentation is not always evident in prior research, which often stops short at model validation without extensive deployment preparation.

In summary, while prior works have laid a solid foundation for predictive maintenance in server infrastructures using ML and DL techniques, this project contributes by (1) integrating more diverse and representative datasets, (2) advancing feature engineering with derived and temporal features, and (3) offering a holistic system that combines

CHAPTER 2

real-time monitoring (HwiNFO, Zabbix) with an explainable, production-ready predictive model. These enhancements directly address limitations identified in the literature and support more resilient and adaptive server maintenance strategies.

CHAPTER 3

Proposed Method/Approach

This chapter outlines the method used to develop a predictive maintenance model for server failure prediction. It covers the system requirements, design, key equations, architecture, challenges, and project timeline, providing a clear framework for building and deploying the system.

3.1 System Development Models

When developing a predictive maintenance system for server failures, it is important to choose the right development model. Different models provide different ways of planning, building, and improving the system. Below are the models considered in this project.

3.1.1 System Development Model 1 – Waterfall Approach

The Waterfall model has a step-by-step approach. Each stage (such as requirements, design, implementation, testing, and deployment) must be completed before moving on to the next one.

- Strengths: Simple, structured, and easy to follow.
- Weaknesses: Not flexible, if something needs to change later, it is very difficult to go back and fix it.

For this project, the Waterfall model is not the best option because predictive maintenance requires testing different datasets, improving models, and making changes along the way. A rigid step-by-step approach would slow down progress.

3.1.2 System Development Model 2 – Spiral Model

The Spiral model works in cycles. A prototype is built, tested, and then improved in the next cycle. This repeats until the system is good enough.

- Strengths: Allows early testing, helps identify risks early, and improves gradually.
- Weaknesses: Can be time-consuming and harder to manage if too many prototypes are built.

In this project, the Spiral model would allow early testing of predictive models and dashboards. However, it still may not provide the speed and flexibility needed to handle very large datasets and frequent adjustments.

3.1.3 System Development Model 3 – Agile Development

The Agile model is more flexible. Work is done in small stages (called iterations), and after each stage, results are reviewed and improved. In a data-driven context, this means experimenting with different datasets, adjusting models, and continuously improving the system.

- Strengths: Very flexible, supports experiments, and produces usable results at each stage.
- Weaknesses: Without careful planning, it can lead to too many changes and poor documentation.

This matches well with this project because the FYP2 workflow itself has six stages (from dataset baselines to real-time monitoring), and each stage is like one Agile iteration.

3.1.4 Selected Model

For this project, the Agile Data-Driven model was chosen. It fits best because:

Stage 1	Models were tested on each dataset (Cloud, Storage, HPC, Azure).
Stage 2	A unified schema was built and adjusted as more data was added.
Stage 3	Models were trained and improved with new techniques.
Stage 4	Evaluation and corrections were done iteratively.
Stage 5	The simulation tool was developed step by step with testing.
Stage 6	Pseudo real-time monitoring was added, extending the system further.

Table 3.1: Selected Model Workflow

This way of working made the project flexible, practical, and able to adapt to problems (like class imbalance or unrealistic predictions) while still moving forward.

3.2 System Requirement

3.2.1 Hardware

A laptop was used as the central development and execution platform for this project. It supported the entire predictive maintenance pipeline, including dataset preprocessing, schema unification, domain-specific and unified model training, and evaluation. The laptop was also utilized to host the Streamlit-based simulation tool and run pseudo real-time monitoring experiments through log replay, ensuring a consistent development and testing environment.



Figure 3.2.1: Specifications of Laptop

Description	Specifications
Model	Lenovo Ideapad Gaming 3
Processor	AMD Ryzen 5 5600H with Radeon Graphics
Operating System	Windows 11 Home Single Language
Graphic	NVIDIA GeForce RTX™ 3050 Laptop GPU
Memory	16GB

Table 3.2.1: Specifications of laptop

3.2.2 Software



Figure 3.2.2: Google Colab

Google Colab was primarily used for dataset preprocessing, schema unification, and baseline model development. With GPU acceleration and Google Drive integration, it enabled efficient handling of large-scale datasets such as Google Cluster and Backblaze, supporting iterative feature engineering and experiment versioning.

*Figure 3.2.3: Kaggle*

Kaggle provided additional computational resources for large-scale unified model training and evaluation, particularly with memory-intensive datasets like Azure (35M+ rows) and the integrated 6.3M unified records. Its high-memory runtime and GPU/TPU access were essential for Model Training and Evaluation.

*Figure 3.2.4: Streamlit*

Streamlit was used to develop the interactive simulation tool. It provided a clean interface for domain-specific predictions (Cloud, HPC, Storage, Unified) with features such as real-time visualization, history tracking, and actionable recommendations. Additionally, `st_autorefresh()` and looping logic were used to enable pseudo real-time monitoring through log replay in Pseudo Real-Time Monitoring.

3.2.3 Data Processing Environment



Figure 3.2.5: Jupyter Notebook

Jupyter Notebook was the main environment for data analysis, feature engineering, and model experimentation. Pandas and NumPy handled preprocessing and schema tasks, while Scikit-learn, XGBoost, and Random Forest supported model training and evaluation. Matplotlib was used for visualizing temporal trends and performance metrics, and Imbalanced-learn (SMOTE, stratified sampling) addressed severe class imbalance. Models were saved and deployed with Joblib and Pickle, ensuring smooth integration with the Streamlit simulation tool. Together, these tools enabled a complete pipeline from raw data processing to pseudo real-time predictive maintenance simulation.

3.3 Functional Requirements

Category	Requirements
User Requirements	<ul style="list-style-type: none"> • A dashboard interface to visualize system health and risk levels. • Real-time monitoring of server performance. • Alerts and notifications when potential failures are detected.
System Requirements	<ul style="list-style-type: none"> • Ability to predict failures across multiple domains (Cloud, Storage, HPC). • Calculation of composite risk scores based on system metrics. • Support for pseudo real-time updates using streaming or batch log simulation.
Performance Requirements	<ul style="list-style-type: none"> • Predictive models should achieve $AUC \geq 0.80$ across datasets. • Optimized memory usage for large datasets (up to tens of millions of records).

	<ul style="list-style-type: none">• Inference speed of less than 1 second per batch to support near real-time monitoring.
--	---

Table 3.2: Functional Requirements

3.4 Timeline




Figure 3.5: Timeline

Current Semester Plan

This semester focused on executing the full predictive maintenance workflow, starting from dataset-specific baselines to a unified, deployment-ready prototype. Public datasets such as Google Cluster, Backblaze, CINECA HPC, and Azure Traces were collected, cleaned, and modeled to establish per-dataset baselines. A 24-column unified schema was then developed to integrate heterogeneous metrics across domains, enabling cross-domain model training. Advanced modeling strategies were applied, including domain-specific models and a unified transfer learning approach, with techniques such as anomaly detection, stratified sampling, and feature importance analysis to manage imbalance and ensure interpretability. A simulation dashboard was built in Streamlit, implementing mathematical risk scoring functions to replace heavy model loading. Finally, pseudo real-time monitoring was introduced using batch log streaming, providing near-live visualization, trend analysis, and proactive alerts. The result is a functional end-to-end PM system prototype that demonstrates predictive maintenance for server failures in virtual environments.

Next Semester Plan

Next semester will emphasize refining and validating the prototype under more realistic monitoring scenarios. This includes extending the pseudo real-time dashboard into a more robust live monitoring tool, integrating with telemetry agents such as HWiNFO and ZabbixAgent for short-term data collection, and enhancing the unified schema with real-time input handling. Additional improvements will focus on optimizing risk models for scalability, strengthening robustness features (e.g., session efficiency, error handling), and incorporating retraining strategies to adapt to evolving workloads. The system will also be tested for cross-domain adaptability, ensuring reliability across cloud, HPC, and storage environments. The semester will conclude with comprehensive documentation, user-friendly simulation demonstrations, and final thesis preparation to showcase a complete predictive maintenance pipeline ready for academic defense and practical extension.

3.5 Estimated Cost

Items	For FYP Development	For Commercialisation
Hardware (Laptop / Server)	Existing / 0	RM4,000+ for high-end
Cloud Services (Kaggle, Google Colab)	Free tier	RM50 per month
Software / Tools (Python, Libraries)	Free	Enterprise Licenses needed
Storage	Free tier	Depending on volume

Table 3.5: Estimated Cost

3.6 Implementation of Issues and Challenges

One of the primary challenges in this project was the limited availability of public, domain-relevant datasets for server predictive maintenance. While datasets such as Google Cluster, Backblaze, CINECA HPC, and Azure Traces provided diverse telemetry, they still lacked certain critical, component-level hardware failure records typically found only in proprietary industry logs. To address this gap, the project

combined these large-scale public datasets with a constructed unified schema and supplemented experimentation using short-term monitoring data collected with tools like HWiNFO and ZabbixAgent. This approach enabled realistic prototyping but inevitably limited the scope compared to long-term, production-scale telemetry.

A second major challenge was the heterogeneity of datasets integrated across cloud workloads, storage devices, and HPC environments. Each dataset differed in structure, sampling rate, and labeling scheme (e.g., EVICT events in Google Cluster vs. SMART failures in Backblaze vs. thermal anomalies in CINECA). Designing the 24-column unified schema required extensive preprocessing, normalization, and careful mapping of domain-specific metrics into standardized features. Ensuring consistency while retaining domain-relevant failure signals was both time-intensive and technically complex.

The feature engineering stage posed additional complexity, particularly in balancing interpretability and predictive power. For example, composite stress indices were designed for cloud workloads, exponential thermal risk curves were used for HPC, and SMART health degradation scores were constructed for storage devices. While these features improved predictive accuracy, they also required careful calibration to avoid deterministic patterns or overfitting, especially given class imbalance across domains.

Model training and evaluation introduced further challenges. Severe imbalance — such as the 0.07% HDD failure rate in Backblaze — made conventional supervised learning unreliable. Although techniques like anomaly detection and stratified sampling improved performance, they sometimes produced inflated AUC scores with limited real-world generalizability. Correcting data leakage (e.g., in early HPC experiments) was also necessary to maintain realistic results.

The simulation and deployment stages highlighted practical integration issues. Building a pseudo real-time dashboard in Streamlit required efficient mathematical formulations to replace heavy model loading, ensuring that live predictions could be generated with minimal system overhead. Memory efficiency, latency reduction, and reproducibility (through session limits and seed-based randomness) were critical to make the tool practical within project constraints.

Another persistent limitation was time and resource constraints. Due to the short project timeline, real-time monitoring could only be demonstrated through log streaming and pseudo real-time simulation rather than full hardware integration. Hardware resource limits (single laptop environment, no dedicated servers) also restricted the scale of experiments, model complexity, and ability to test deployment in production-grade environments.

Despite these challenges, the project delivered a novel, end-to-end predictive maintenance workflow: from per-dataset baselines to a unified schema, advanced domain-aware modeling, and a professional simulation dashboard. Its key contributions lie in integrating heterogeneous datasets, resolving deterministic and anomaly dominance issues, and creating a deployment-ready monitoring tool that balances accuracy, interpretability, and resource efficiency — providing a foundation for future extensions in real-world server infrastructures.

CHAPTER 4

System Design

This section provides a high-level overview of the predictive maintenance application's key features and structure. The goal is to offer a clear understanding of the system's functionality and how the various components interact within the overall design.

4.1 System Architecture Diagram

The following is the system architecture, where the developed solution consists of multiple stages corresponding to dataset handling, model training, and simulation tool deployment. Each stage has its own flow, but they are integrated into a single pipeline that enables predictive maintenance for cloud, HPC, and storage domains.

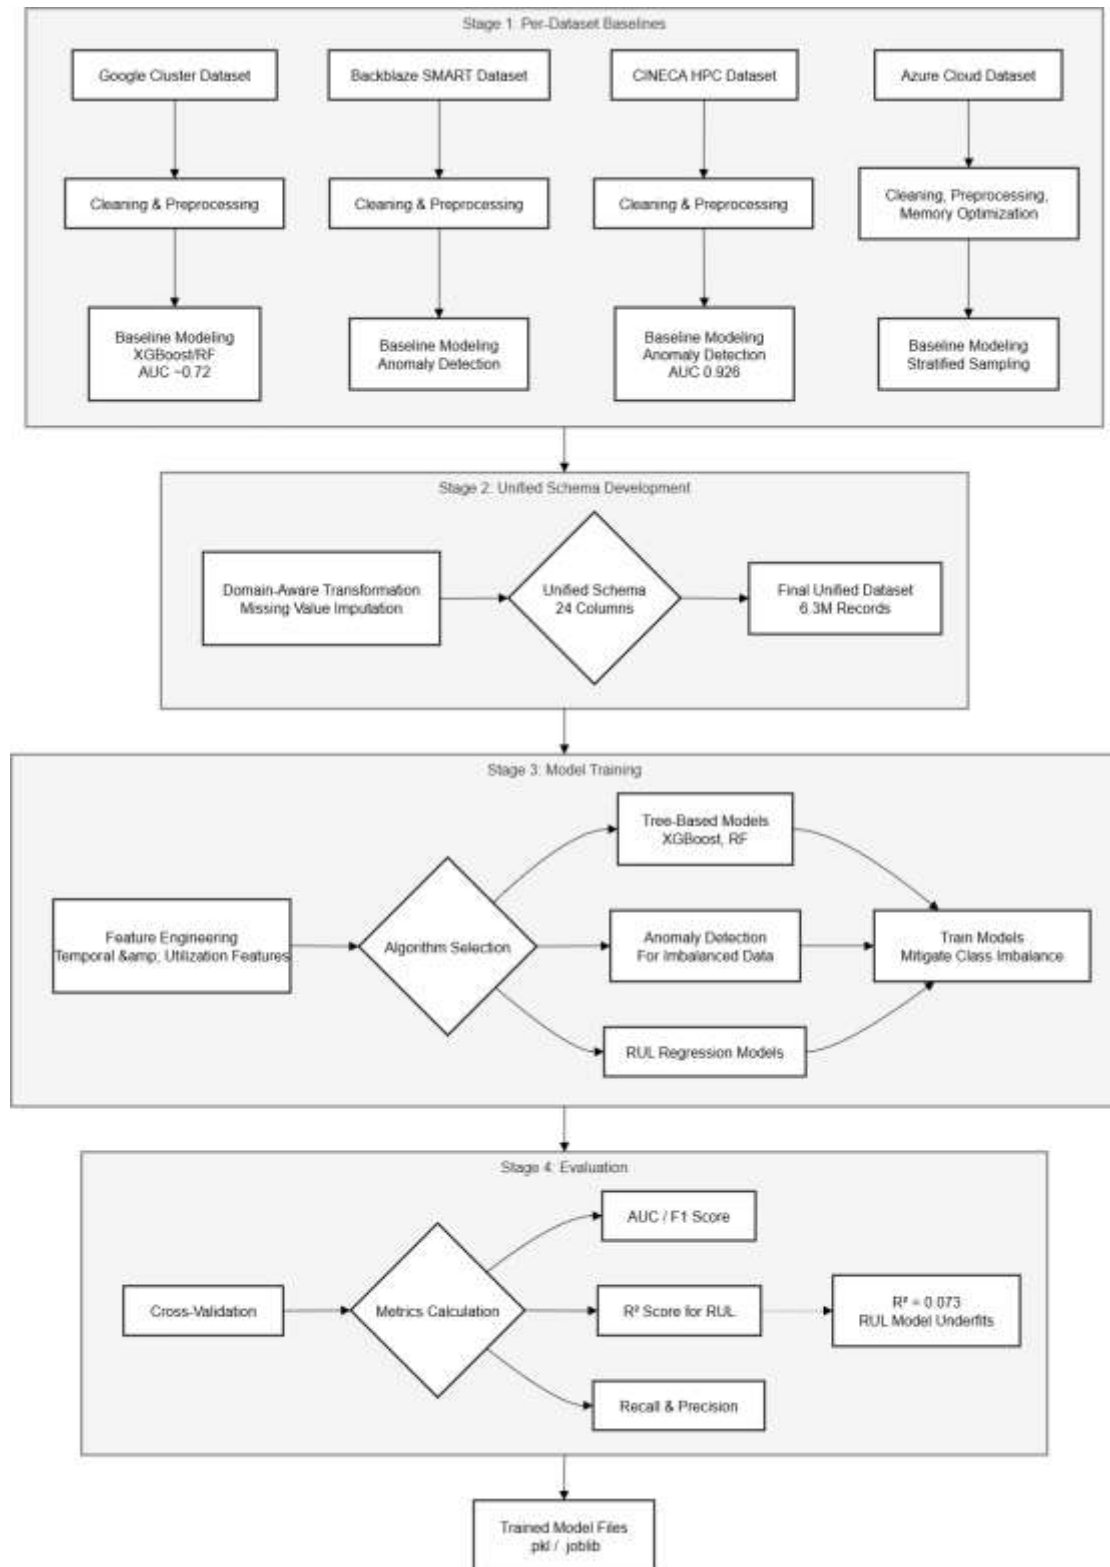


Figure 4.1.1: Data Unification and Model Training Pipeline

Figure 4.1.1 illustrates the initial, foundational stages of the predictive maintenance system, which operate in a batch processing manner. The workflow begins with Stage 1, where four distinct datasets (Google Cluster, Backblaze, CINECA HPC, and Azure

Cloud) are processed independently. Each dataset undergoes its own specific cleaning and preprocessing routines, acknowledging their unique formats and challenges, such as Azure's large scale requiring memory optimization. Baseline models are then trained on these individual datasets, employing algorithms suited to their specific characteristics—like XGBoost for Google Cluster and specialized anomaly detection for the severely imbalanced Backblaze and complex CINECA HPC data. This stage establishes performance benchmarks (e.g., AUC scores) for each domain.

The processed data from all domains then flows into Stage 2, Unified Schema Development. Here, a critical step of domain-aware transformation is applied to convert the heterogeneous data into a consistent, 24-column schema. This involves intelligent handling of missing values and ensures all data speaks a "common language," resulting in a final, unified dataset of 6.3 million records. Stage 3, Model Training, takes this unified data and performs feature engineering to create predictive features based on time and resource utilization. Multiple model types are trained, including tree-based classifiers, anomaly detection algorithms, and Remaining Useful Life (RUL) regression models, all while employing techniques to mitigate class imbalance. Finally, Stage 4, Evaluation, rigorously validates these models using cross-validation and a suite of metrics. This stage confirmed the strong performance of the classification models, identified and corrected issues like data leakage, and highlighted the inherent difficulty of predicting RUL with the available data, as evidenced by the low R^2 score. The final output of this entire pipeline is a set of trained and validated model files, ready for deployment.

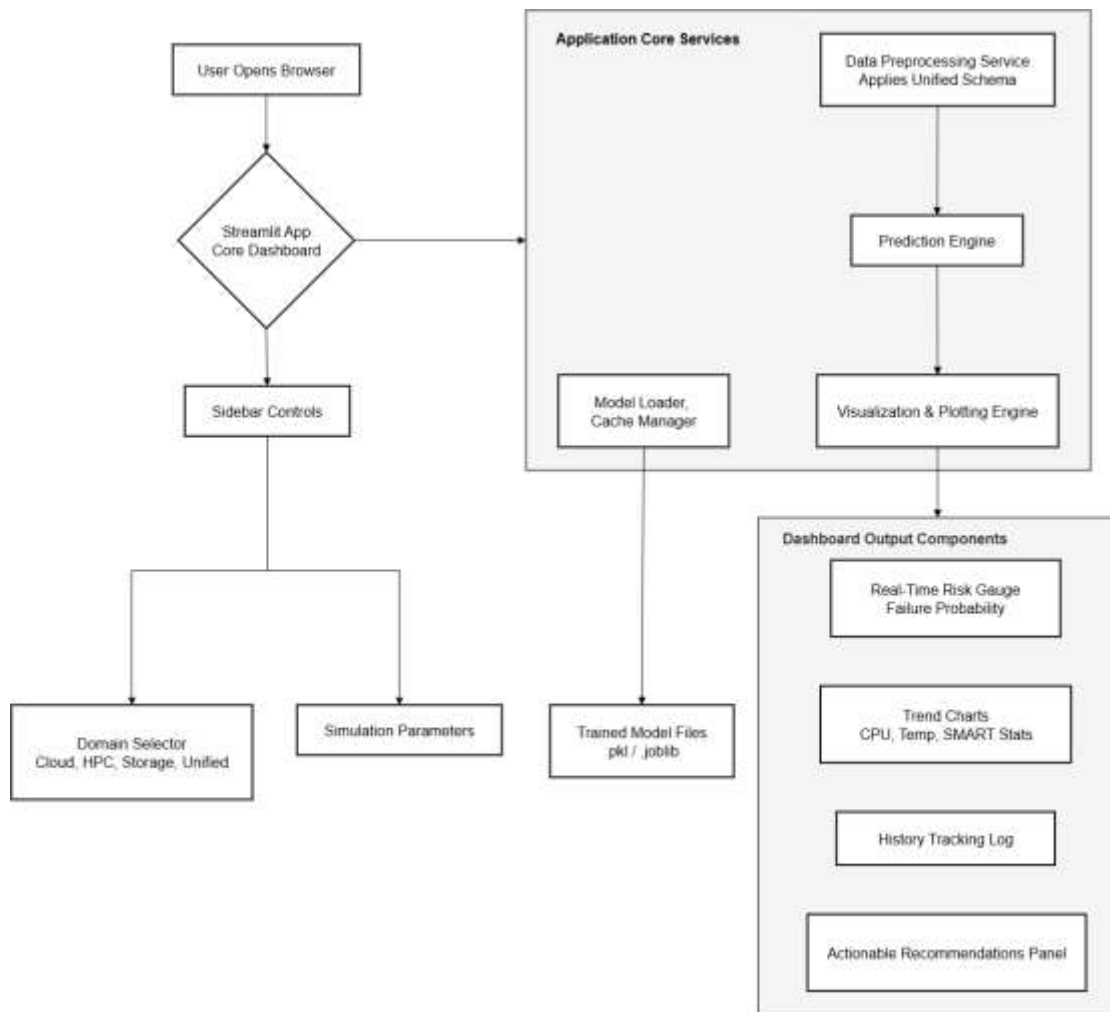


Figure 4.1.2: Simulation Tool Architecture

Figure 4.1.2 depicts the architecture of the Streamlit-based simulation tool, which serves as the user-facing application for interacting with the trained models. The system is centered around a Core Dashboard that a user accesses via a web browser. User interaction begins in the Sidebar Controls, where they can select the specific domain they wish to simulate (e.g., Cloud, HPC) and adjust various simulation parameters. This user input drives the behavior of the application's core services.

The heart of the application consists of four integrated services. First, the Model Loader fetches the appropriate pre-trained model files (e.g., .pkl) generated from the previous pipeline. Second, the Data Preprocessing Service ensures that any input data is automatically transformed to match the 24-column unified schema created in Stage 2, maintaining consistency between training and prediction. Third, the Prediction Engine runs this preprocessed data through the loaded model to generate predictions. Fourth, the Visualization Engine takes these results and dynamically renders them into

the Dashboard Output Components. These components include a real-time risk gauge showing failure probability, interactive trend charts for key metrics, a history log for tracking events, and a panel providing actionable recommendations, effectively translating model outputs into operational insights for the user.

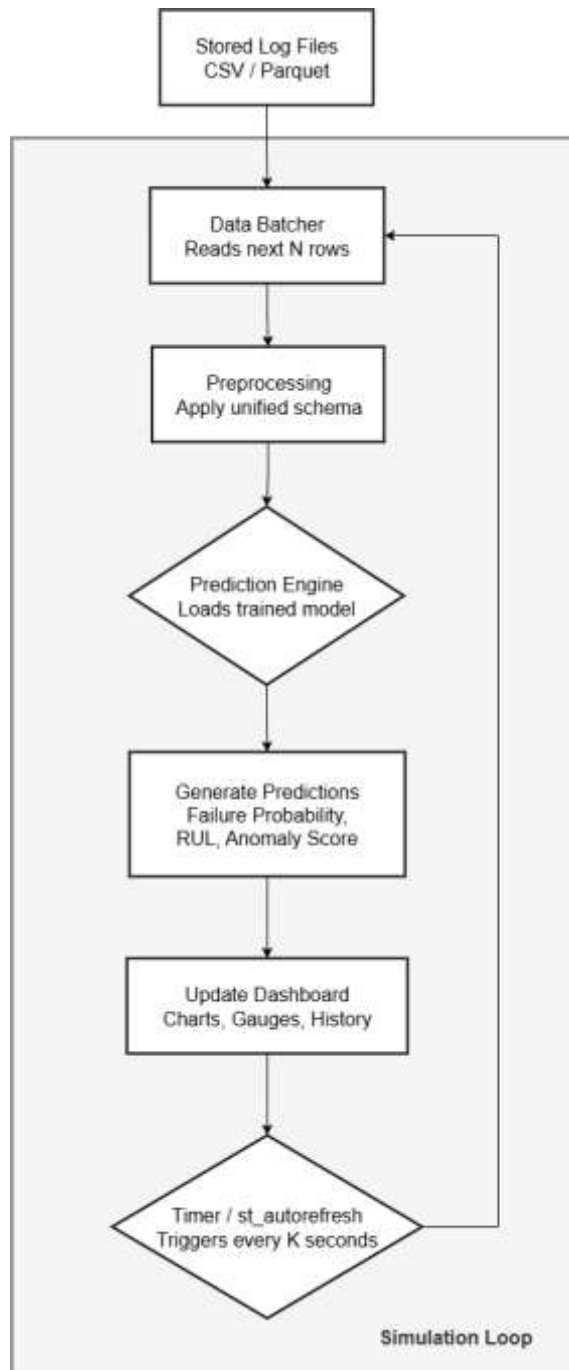


Figure 4.1.3: Pseudo Real-Time Monitoring Simulation

Figure 4.1.3 details the architecture of the pseudo real-time monitoring system that simulates a live production environment. Since the project operates in a virtual environment without direct access to live server telemetry, this stage ingeniously mimics a live data stream. The process is built on a timed Simulation Loop, triggered at regular intervals (e.g., every few seconds) by a timer function within the Streamlit app.

The loop starts by reading a small, new batch of data from the Stored Log Files (CSV/Parquet) that were used in training. This Data Batcher component acts as a stand-in for a live data streamer. Each new batch of data is then passed through the same Preprocessing and Prediction Engine used in the simulation tool (from Figure 2). The generated predictions (failure probabilities, etc.) are immediately sent to update the dashboard visualizations, creating the impression of a live, continuously updating monitoring system. Once the update is complete, the loop waits for the next timer tick to begin the process again with the next batch of data. This architecture demonstrates the complete operational flow of a predictive maintenance system—from data ingestion to actionable insight—in a way that is both effective for demonstration and feasible for a project environment.

4.2 Use Case Diagram

This system provides a comprehensive solution for predicting server failures across diverse computing environments (Cloud, HPC, Storage) by leveraging machine learning on operational data. The workflow is designed around four primary actors, each interacting with a specific subsystem of the platform to achieve a complete pipeline from data preparation to real-time monitoring and user interaction.

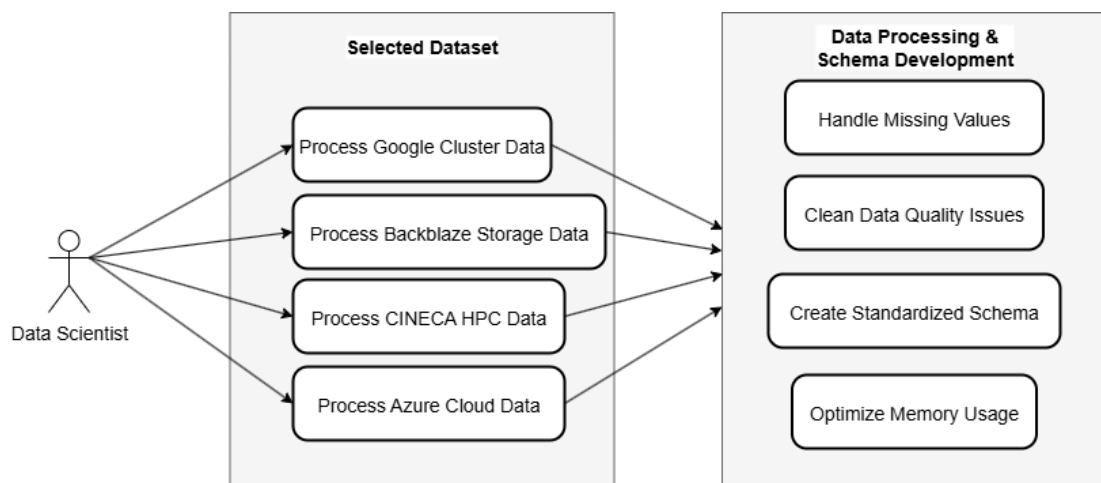


Figure 4.2.1: Data Scientist Use Cases

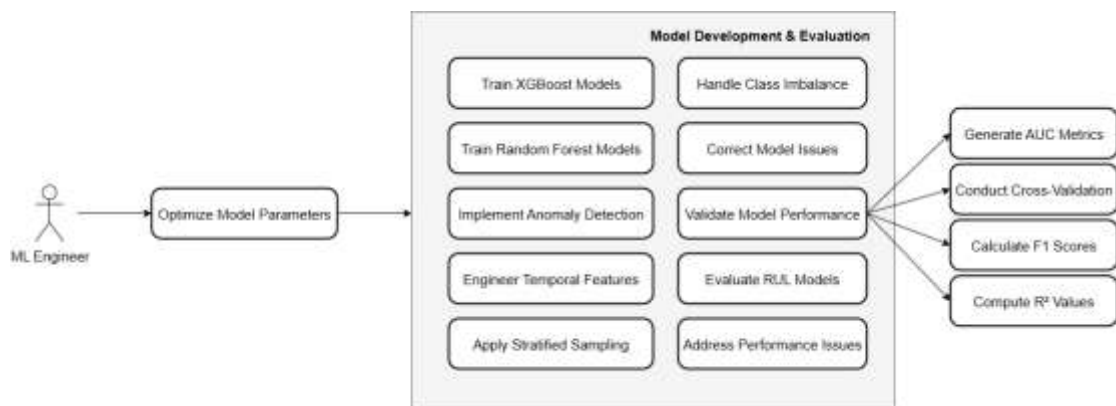


Figure 4.2.2: ML Engineer Use Cases

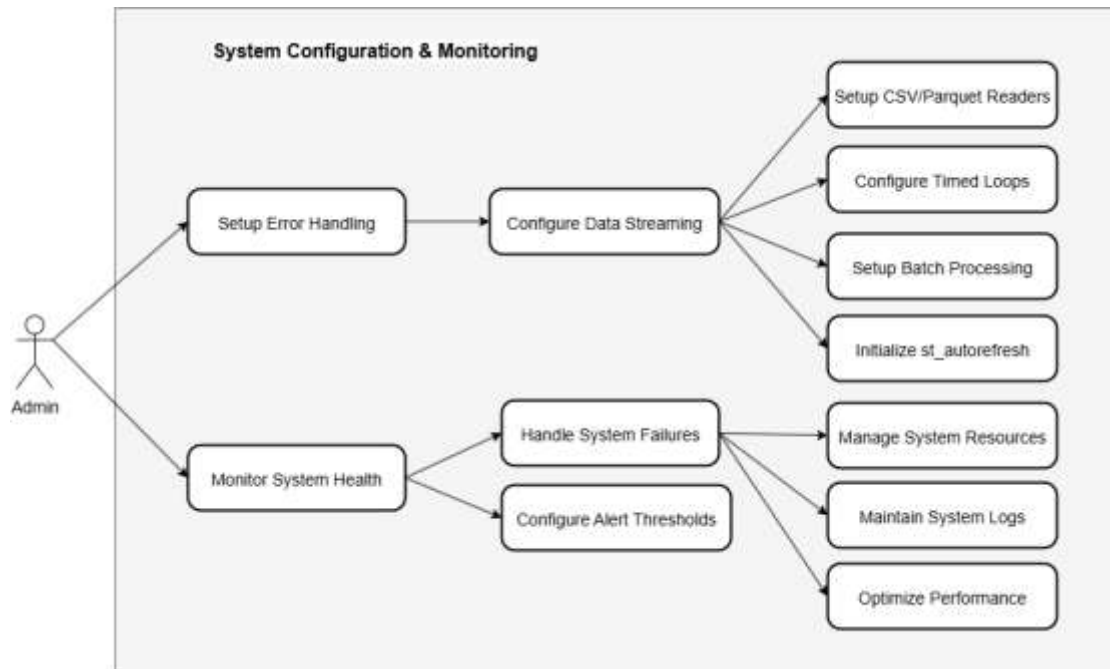


Figure 4.2.3: System Admin Use Cases

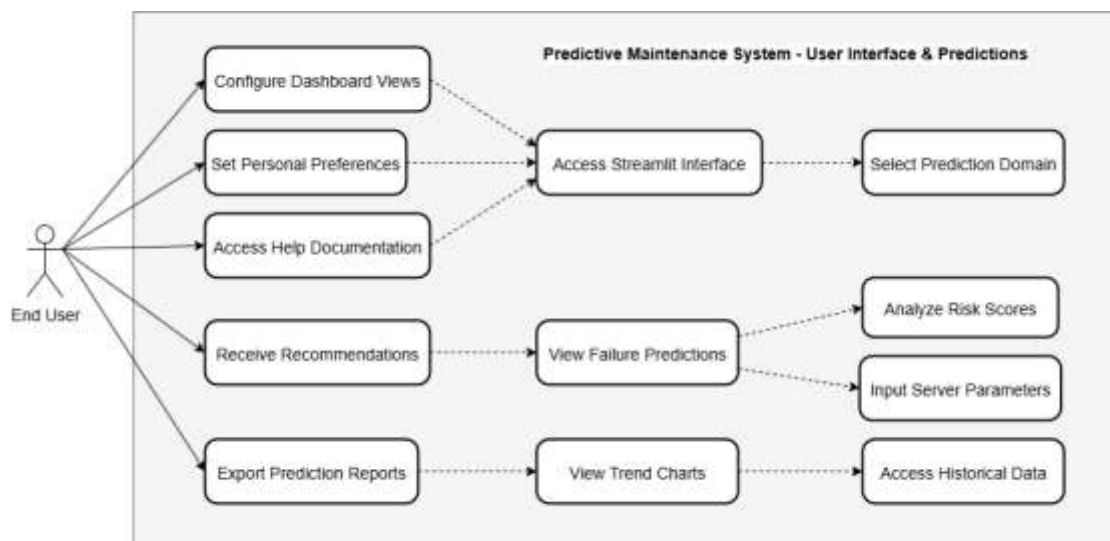


Figure 4.2.4: End User Use Cases

Actor	Stage(s)	Main Role	Outcome
Data Scientist	1-2	Clean and unify datasets into a standardized schema.	Reliable unified dataset for training.

ML Engineer	3-4	Build, tune, and validate predictive models.	Robust models ready for deployment.
End User	5	Use dashboard to view predictions and insights.	Clear server health insights for proactive action.
Admin	6	Configure pseudo real-time monitoring and system health.	Stable, responsive monitoring platform.

Table 4.2: Actor Interactions

4.3 Activity Diagram

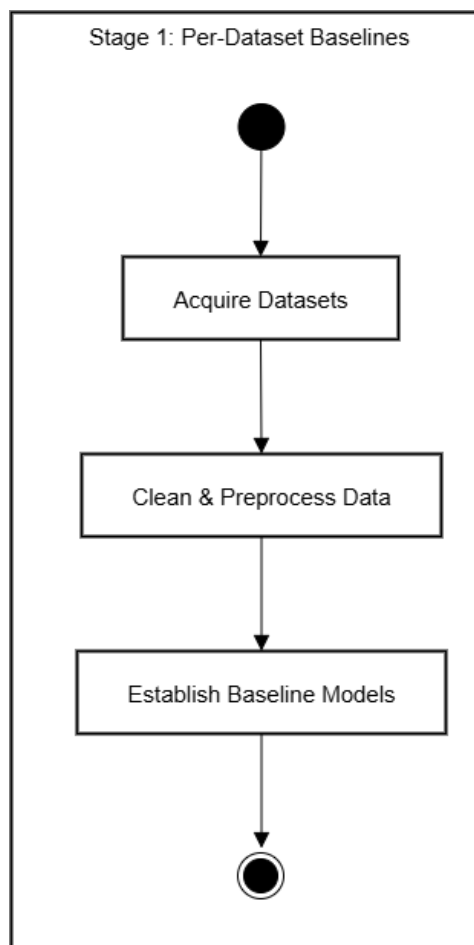


Figure 4.3.1: Per-Dataset Baselines

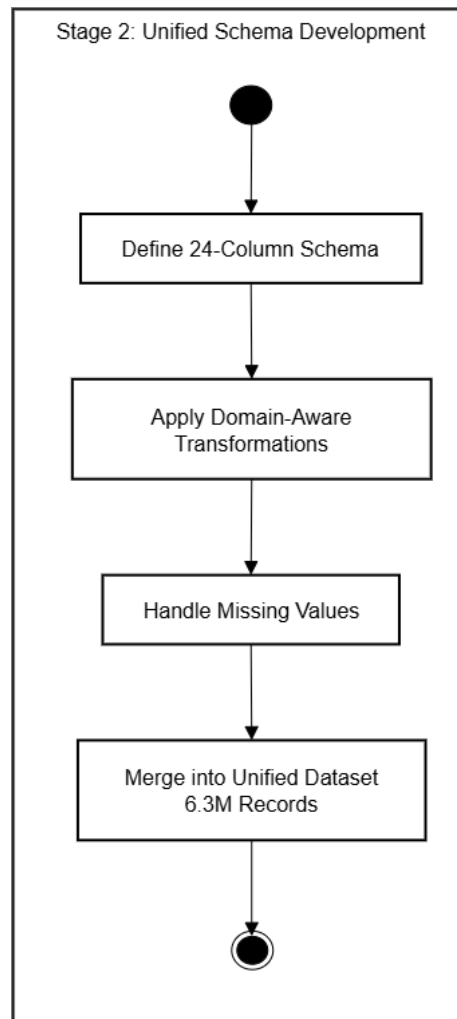


Figure 4.3.2: Unified Schema Development

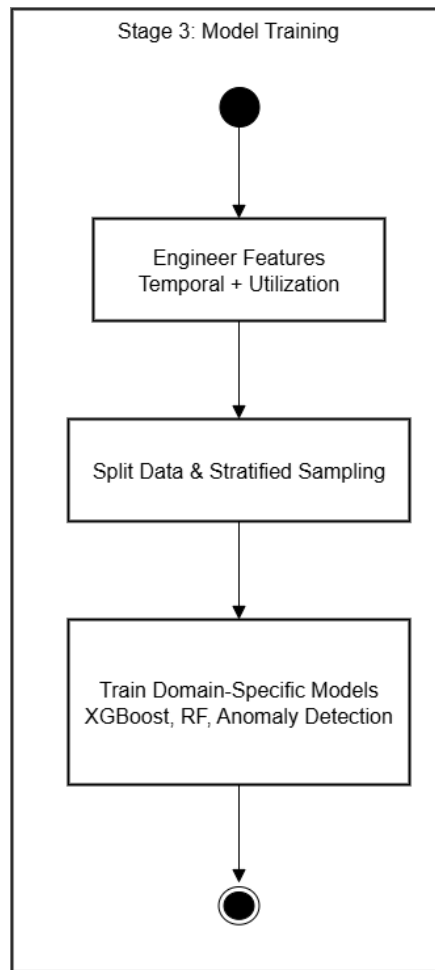


Figure 4.3.3: Model Training

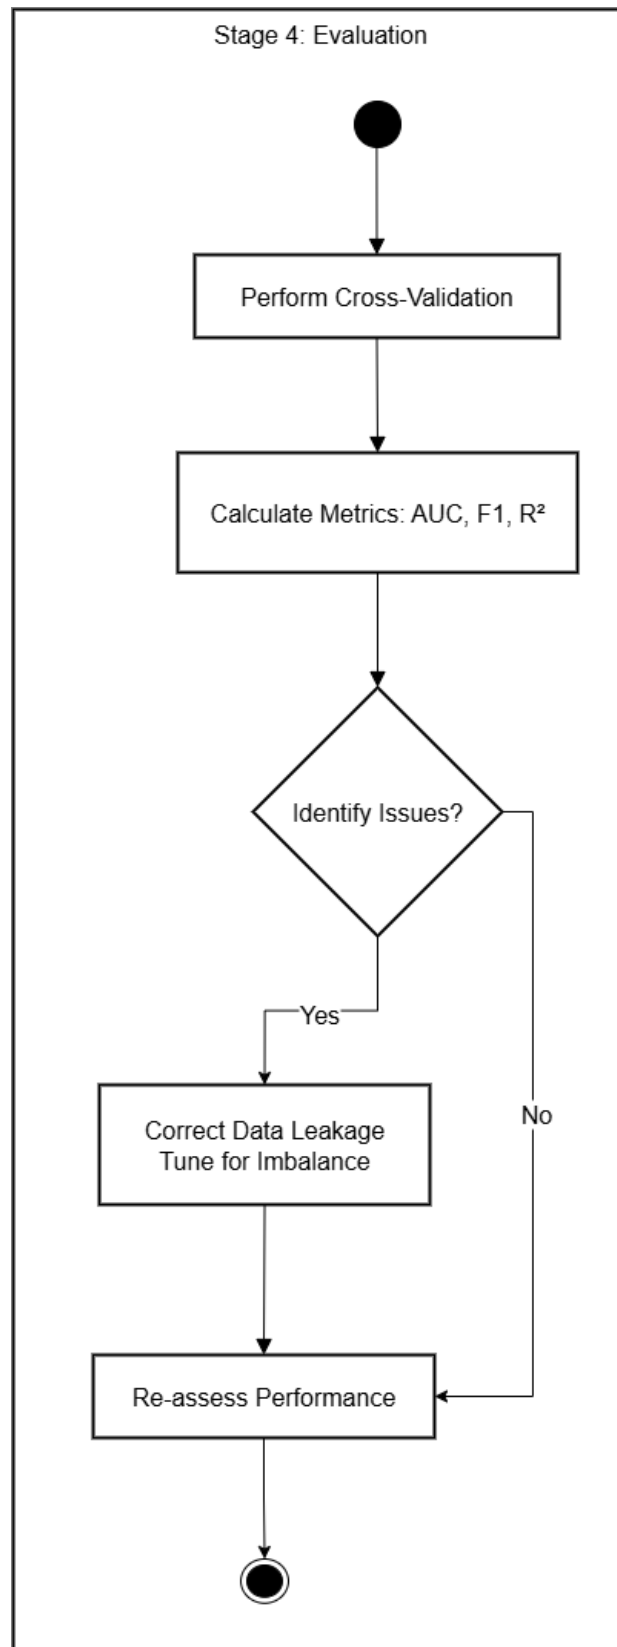


Figure 4.3.4: Model Evaluation

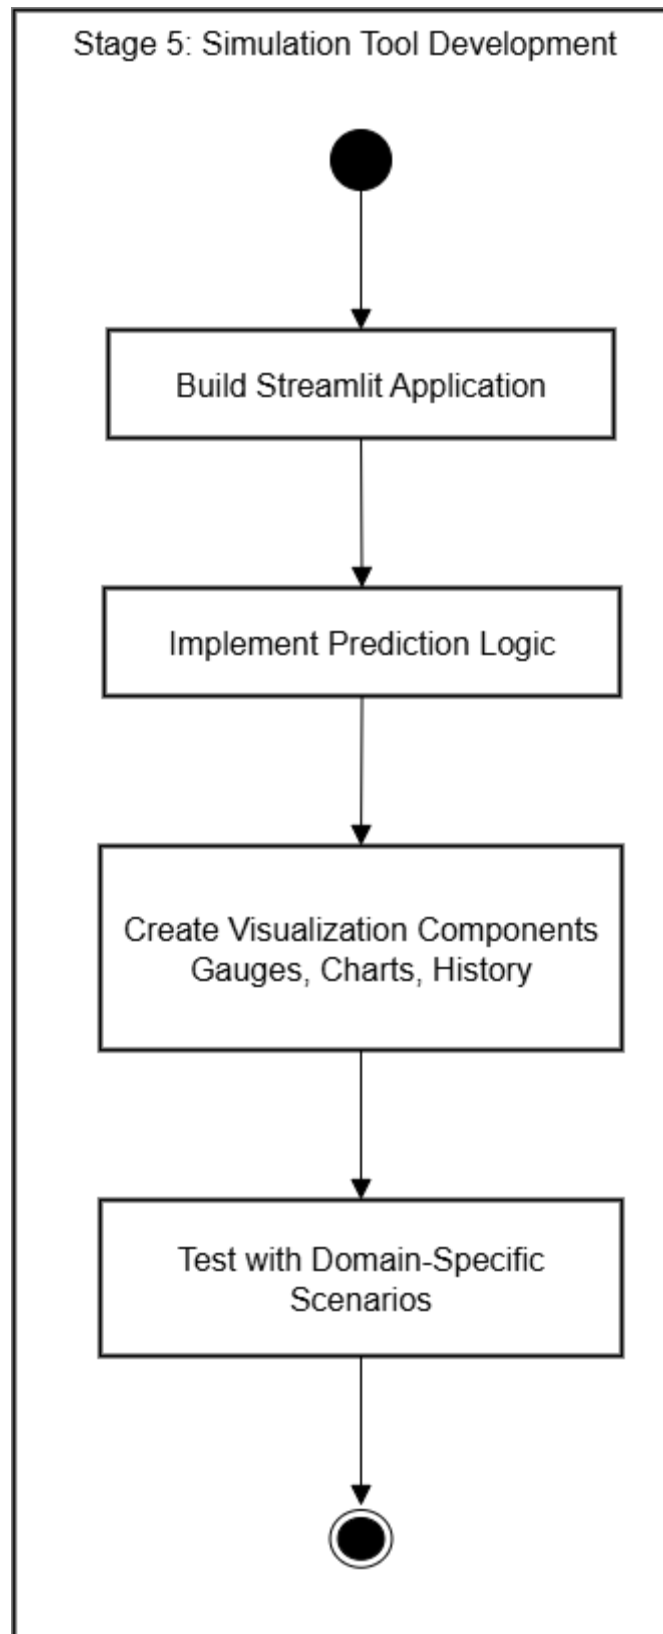


Figure 4.3.5: Simulation Tool Development

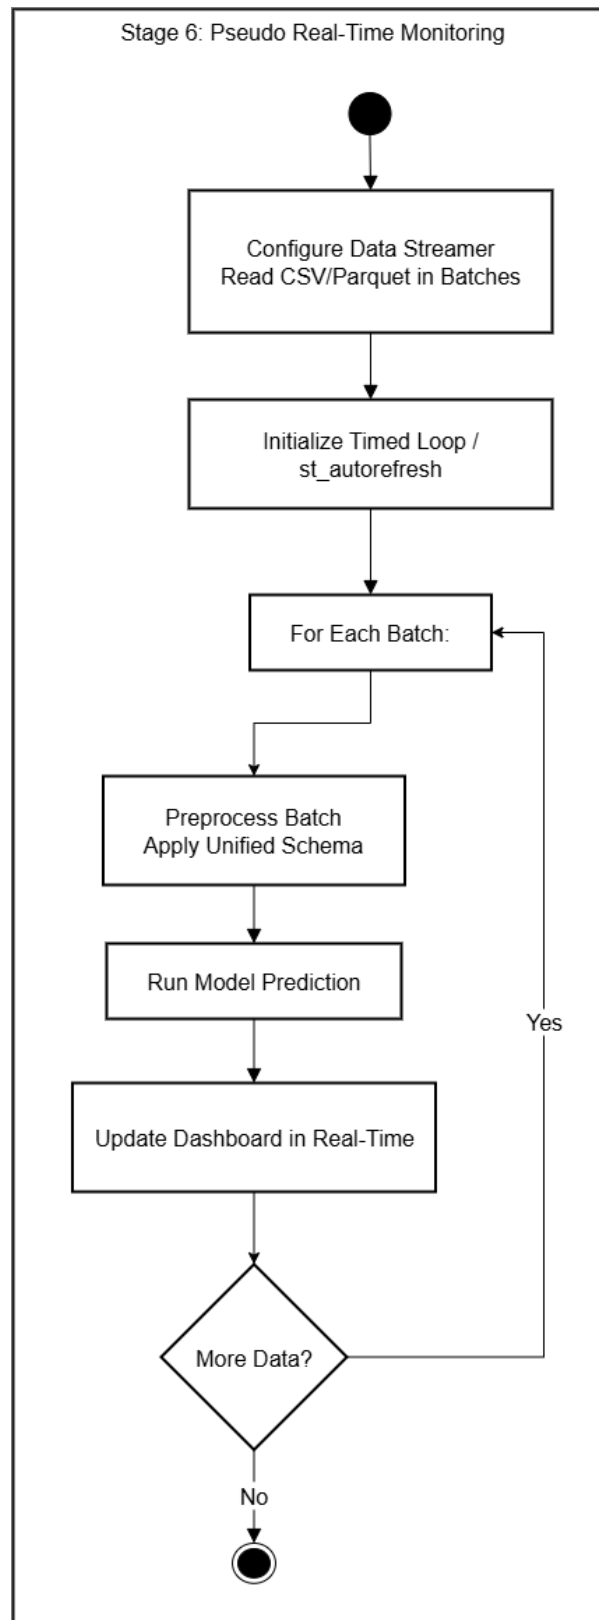


Figure 4.3.6: Pseudo Real-Time Monitoring

The activity diagram illustrates the six stages of the predictive maintenance workflow.

Stage 1 (Per-Dataset Baselines) begins with acquiring datasets, followed by cleaning,

preprocessing, and building baseline models for each domain. Stage 2 (Unified Schema Development) focuses on defining a 24-column schema, applying domain-aware transformations, handling missing values, and merging all sources into a unified dataset of 6.3M records. Stage 3 (Model Training) covers feature engineering, stratified sampling, and training domain-specific models such as XGBoost, Random Forest, and anomaly detection. Stage 4 (Evaluation) validates model performance through cross-validation and metric calculation (AUC, F1, R^2), with decisions to correct issues like data leakage or class imbalance before finalizing results. Stage 5 (Simulation Tool Development) involves building a Streamlit application, implementing prediction logic, designing visualization components, and testing the tool with domain-specific scenarios. Finally, Stage 6 (Pseudo Real-Time Monitoring) configures batch-based data streaming, initializes timed loops, preprocesses incoming data, generates predictions, updates the dashboard, and repeats the process until all data is processed, completing the workflow pipeline.

4.4 System Block Diagram

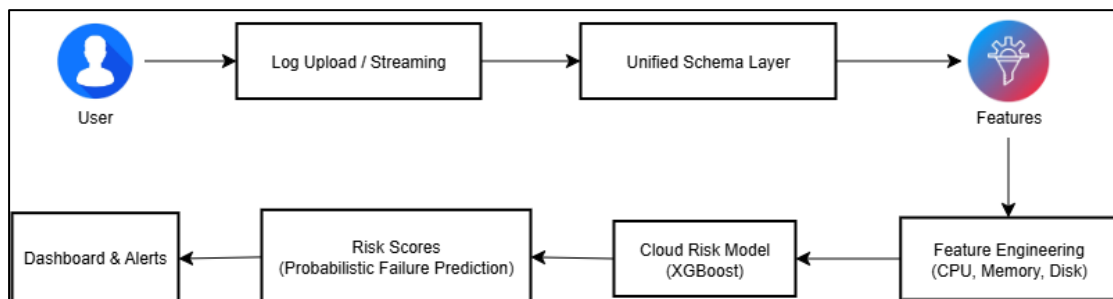


Figure 4.4.1 – Cloud Failure Prediction Service

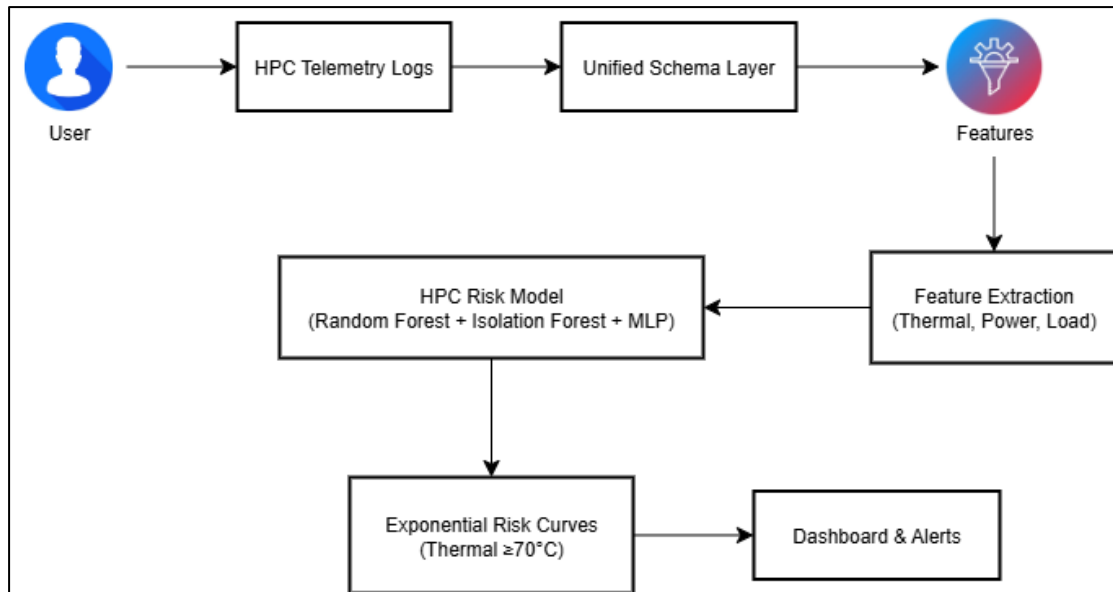


Figure 4.4.2 – HPC Thermal-Power Risk Service

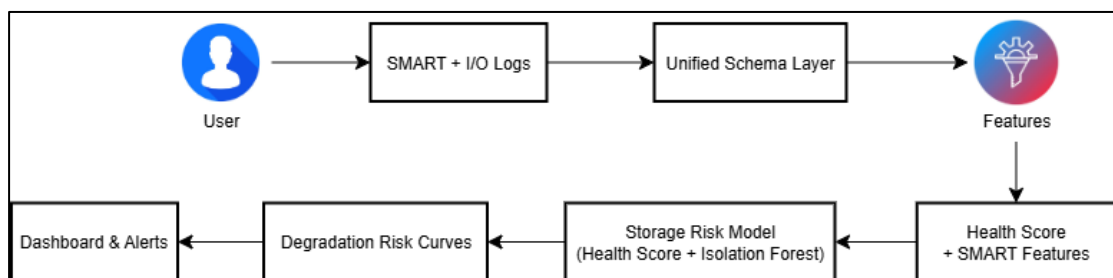


Figure 4.4.3 – Storage Health Monitoring Service

4.4.1 User

The user plays a central role in initiating the predictive maintenance workflow. To begin the process, the user either uploads historical logs or streams operational data into the system. The design of the application emphasizes flexibility, ensuring that services can be activated manually or operated in simulation mode. This dual approach allows for both batch analyses of large datasets and continuous pseudo real-time monitoring. For example, a user may enable the cloud workload stress predictor through the dashboard to examine risk levels in virtual machine traces.

4.4.2 Virtualized Environment (Cloud, HPC, and Storage Systems)

The system has been designed to support a diverse range of virtualized environments, spanning cloud workloads, high-performance computing (HPC) systems, and storage

infrastructures. In cloud domains such as Google and Azure, the primary focus is on monitoring CPU, memory, and disk utilization to identify workload-induced stress. Within HPC environments like CINECA's M100 supercomputer, thermal and power-based anomalies are prioritized to ensure reliability during intensive computations. Storage systems, particularly those modeled from Backblaze HDD datasets, rely on SMART attributes and health degradation patterns to capture early signs of hardware failure. All of these environments are processed through a unified schema containing twenty-four standardized columns, ensuring interoperability and comparability across domains.

4.4.3 Risk Modeling Engine

At the core of the system lies the risk modeling engine, which is responsible for transforming raw operational metrics into probabilistic failure predictions. The engine employs three domain-specific models tailored to the characteristics of each environment. In the cloud domain, a composite stress model integrates CPU (40%), memory (35%), and disk (25%) utilization to estimate risk. In the HPC domain, exponential risk curves centered on thermal thresholds ($\geq 70^{\circ}\text{C}$) provide a more accurate reflection of overheating conditions, complemented by power consumption factors. The storage domain employs health score degradation as its primary indicator, supplemented by anomaly detection fallbacks for rare events. Beyond these domain-specific models, a unified transfer-learning framework applies cross-domain adjustment factors, enabling more generalized predictions that remain applicable across heterogeneous environments.

4.4.4 Simulation and Dashboard Service

To make predictive maintenance insights accessible, the system incorporates a simulation and dashboard service built with Streamlit. This service provides a professional and user-friendly interface for risk visualization and monitoring. Through the dashboard, users can interact with gauges that reflect failure probabilities relative to critical thresholds, while historical trend analysis supports long-term performance evaluation. The system further allows users to replay logs in real time, using timed refreshes to simulate continuous monitoring scenarios. Color-coded alerts and

actionable recommendations are presented to ensure that critical risks are not only detected but also translated into clear and understandable decisions for operators.

4.4.5 Data Services

Supporting the predictive workflow is a robust data services layer designed to process and standardize heterogeneous datasets. The unified schema acts as the foundation, transforming raw inputs into a consistent and comparable format. On top of this schema, anomaly detection modules are deployed, with models such as Isolation Forest, Random Forest, and XGBoost tuned to each domain. Mathematical risk curves replace simplistic threshold-based triggers, introducing non-linear scaling that prevents deterministic or unrealistic classifications. This approach ensures that predictions remain both accurate and probabilistic, avoiding rigid binary outcomes.

4.5 System Components

Component	Description	Key Functions	Tools/Technologies
Data Pipeline	Handles ingestion and preprocessing of large-scale logs.	Cleaning, anomaly removal, chunk-based processing, feature extraction.	Python (Pandas, Dask), Parquet/CSV.
Unified Schema	Standardized schema with 24 attributes across domains.	Aligns temporal, workload, SMART, and thermal metrics.	Custom schema design, Pandas, SQL-like joins.
Modeling Engine	Implements ML and anomaly models.	XGBoost (Cloud/HPC), Isolation Forest (Storage), Transfer learning (Unified).	Scikit-learn, XGBoost, PyTorch (MLP).

Simulation Tool	Lightweight mathematical risk prediction system.	Domain-specific risk scoring, probability calibration, parameter adjustment.	Python functions, composite equations.
Monitoring Dashboard	User interface for interactive monitoring.	Risk gauges, trend analysis, alerts, recommendations.	Streamlit, Matplotlib/Plotly.

Table 4.5: Algorithm of Core Pipeline

4.6 Dataset Summary Table

Dataset Name	Source	Size	Type	Relevance
Google Cluster Data v2019 [34]	Google Research	~40 GB	Real	Large-scale server workload and failure logs, used for per-dataset baseline (AUC ~0.72).
Backblaze Hard Drive Data Q2 2025 [32]	Backblaze Inc.	~15 GB	Real	HDD S.M.A.R.T. statistics with extremely low failure rate (0.07%), addressed via anomaly detection.
CINECA M100 HPC Data [33]	CINECA Supercomputing Center	~5 GB	Real	Thermal and power consumption data from HPC systems, modeled with anomaly detection (AUC 0.926).

Azure Trace for Packing 2020 [35]	Microsoft (OSDI 2020, Protean Paper)	~50 MB	Real (Cloud Workload Trace)	Captures VM allocation, resource requests, and scheduling behavior in Azure; useful for analyzing packing algorithms, workload patterns, and preemptive eviction in large-scale cloud environments.
Unified Dataset	Constructed (Stage 2)	~6.3M records	Hybrid (Real + Simulated)	24-column schema integrating cloud, HPC, and storage telemetry for cross-domain model training and deployment.

Table 4.6: Summary of Datasets Used

The project relied on five key datasets aligned with the six-stage workflow. The Google Cluster Data was cleaned and modeled with XGBoost and Random Forest, achieving ~0.72 AUC and serving as a baseline for large-scale server workloads. The Backblaze Hard Drive Data provided S.M.A.R.T. metrics with an extremely low failure rate (0.07%), requiring anomaly detection to address class imbalance. The CINECA HPC Data captured thermal and power telemetry from supercomputers, where anomaly detection produced strong results (AUC 0.926). The Azure Trace for Packing 2020 offered insights into VM allocation and scheduling in cloud systems, used to evaluate workload distribution and predictive modeling under resource constraints. Finally, these domains were combined into a Unified Dataset of 6.3M records via a 24-column schema, enabling cross-domain model training, evaluation, and deployment within the simulation and pseudo real-time monitoring stages.

4.7 Sensor Metric to Hardware Component Mapping Diagram

Sensor Metric	Hardware Component	Domain
CPU Utilization & Temperature	CPU / Processor	Cloud & HPC primary driver of resource stress and thermal risk
Memory Usage & Stress Ratios	RAM / System Memory	Cloud & HPC part of composite resource stress modeling
Disk SMART Attributes (Read/Write Errors, Reallocated Sectors, I/O Stats)	Storage (HDD/SSD)	Backblaze & Storage domain – used in health score and anomaly detection
Power Consumption & Voltage	PSU / Node Power	HPC – monitors thermal-power stress, feeds into exponential thermal risk model
Thermal Sensor Readings (Node/Server Temp)	Cooling / Environmental Systems	HPC – drives thermal-dominant risk model
VM Allocation, Eviction, and Lifetime Events	Virtualization Layer / Cloud Scheduler	Google Cluster & Azure – indicates workload stress, early termination patterns, and preemptive failure risk
Composite Risk Scores	Aggregated System Health	Unified domain – combines metrics across CPU, memory, storage, thermal, and workload to provide

		overall	failure
		probability	

Table 4.7: Mapping of Sensor Metrics to Hardware Component

4.8 Data Flowchart

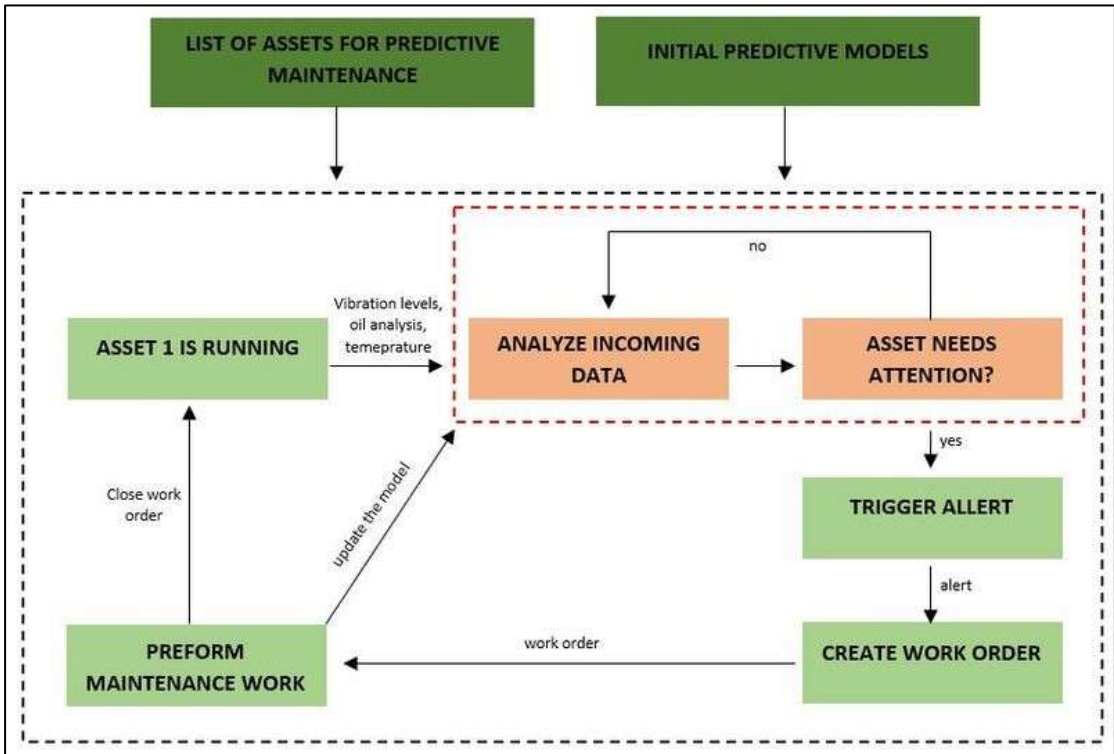


Figure 4.8: Predictive Maintenance Workflow

The predictive maintenance workflow in this project begins with asset data sources from cloud, HPC, and storage domains, which are processed through the unified schema to standardize raw logs into a consistent 24-column format. These preprocessed datasets are then analyzed by both domain-specific and unified predictive models to assess risk levels. The workflow reaches a critical decision point where the system determines whether an asset requires attention. If the models indicate elevated risk, the system triggers alerts through the simulation dashboard, providing recommended maintenance actions. If no risk is detected, the monitoring process continues with incoming data streams. This loop is integrated into the Streamlit-based monitoring tool, which supports pseudo real-time visualization, trend analysis, and proactive alerting, ensuring a continuous and effective predictive maintenance cycle.

4.9 System Equation

Feature/Function	Equation	Description
CPU Failure Probability	Eq. (1)	Estimates CPU failure likelihood based on high temperature ($>75^{\circ}\text{C}$), high utilization ($>90\%$), and long uptime (>5000 h).
Memory Failure Probability	Eq. (2)	Calculates memory failure risk based on high RAM usage ($>85\%$) and abnormal voltage (outside 11.5V – 12.5V).
Mechanical Power	Eq. (3)	Computes mechanical power from torque and rotational speed to capture load conditions relevant to stress.
Temperature-Voltage Ratio	Eq. (4)	Indicates thermal stress relative to voltage stability — useful as a stress feature.
Rolling Mean ($n = 5$)	Eq. (5)	Smooths recent sensor values over 5 -time steps to capture short-term trends and improve robustness.
Rolling Standard Deviation ($n = 5$)	Eq. (6)	Measures short-term variability in sensor values (window = 5). High variation may indicate instability.
Risk Classification Function	Eq. (7)	Converts predicted failure probability into qualitative risk levels for easier interpretation.
Model Training Objective (Random Forest)	Eq. (8)	Minimizes cross-entropy loss with regularization to balance

		fit and complexity in the Random Forest model.
Threshold Optimization (Rolling Mean)	Eq. (9)	Dynamically adjusts decision thresholds by computing a moving average over recent values ($n = 5$).
Final Prediction	Eq. (10)	Binarizes failure probability into final prediction (1 = failure, 0 = no failure) based on decision threshold τ^* .
Tool Wear–Torque Interaction	Eq. (11)	Captures interaction between tool wear (minutes) and torque (Nm) to model compound stress effects.

Table 4.9: Equation used for interaction features of model testing

4.10 Functional Modules in the System

Module	Name	Description
1	Data Ingestion & Preprocessing	Load multi-domain datasets, clean, and standardize into unified schema.
2	Feature Engineering	Compute resource stress, thermal thresholds, SMART-based health indicators.
3	Model Training & Evaluation	Per-domain models (XGBoost, Random Forest, Isolation Forest, MLP). Unified cross-domain model (transfer learning, stratified sampling).
4	Simulation Tool	Composite risk scoring with mathematical calibration. Domain-specific probability bounds.

5	Dashboard & User Interface	Interactive gauges, historical trends, risk alerts, parameter tuning.
6	Pseudo Real-Time Monitoring	Log streaming, periodic updates, continuous risk visualization.

Table 4.10: Functional Modules

4.11 Data & Model Specification Summary

Domain	Dataset Processed	Failure Labels	Final Model	Performance
Cloud (Google Cluster)	1.55M VM logs	EVICT, FAIL, LOST, KILL	XGBoost	AUC 0.8283
Storage (Backblaze)	1.5M HDD SMART logs	HDD Failures	Isolation Forest + Health Score	AUC 0.9272 (F1 = 0.52)
HPC (CINECA M100)	1.25M Thermal/power logs	Thermal/Power anomalies	Fusion: Random Forest + IF + MLP	AUC 0.9445
Cloud Diversity (Azure VM)	2M VM workload traces	Early Termination	XGBoost	AUC 0.9121
Unified	6.3M standardized records	Cross-domain	Transfer Learning Model	AUC 0.8165

Table 4.11: Data & Model Specification

4.12 GUI Design

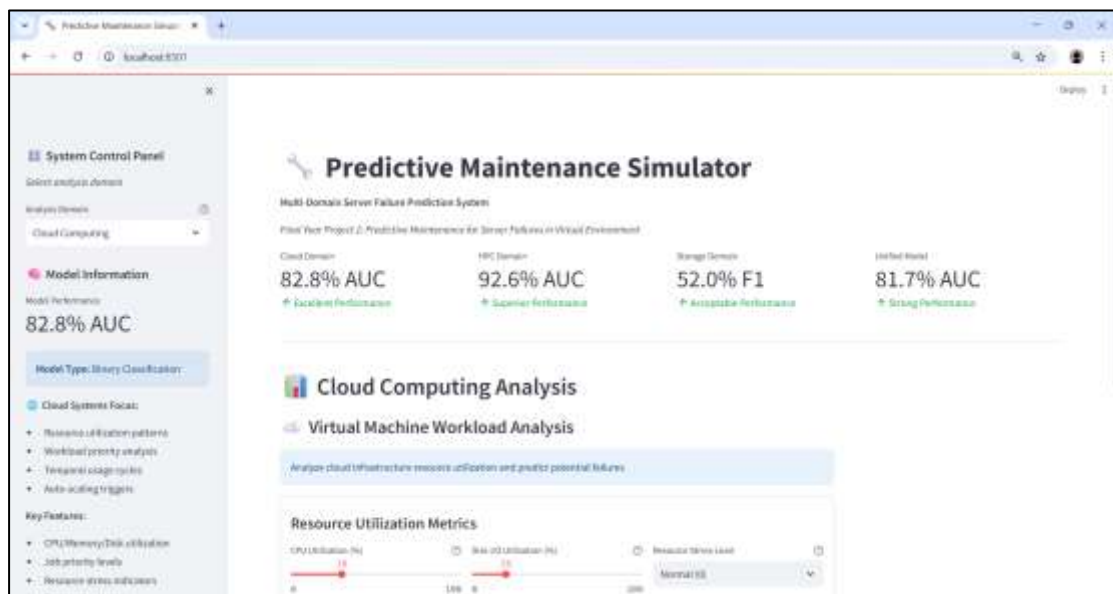


Figure 4.12: GUI Design

The graphical user interface of the predictive maintenance system is implemented as a Streamlit dashboard that closely follows the six-stage FYP2 workflow. The dashboard provides a clean and professional layout designed to prioritize actionable insights over raw data. Users can select the domain of interest—Cloud, HPC, Storage, or Unified—through a dedicated domain selection panel. Real-time risk gauges display the current predicted failure risk using color-coded indicators, while threshold alerts provide immediate visual warnings when critical levels are reached. Trend analysis charts allow users to review historical risk patterns and observe how parameter changes affect predictions over time, supporting operational planning and pattern recognition.

A prediction history and logging feature stores results for each batch, enabling post-analysis and ensuring reproducibility through seed-based randomization. Users can also interact with the dashboard via a parameter adjustment panel, modifying stress multipliers, thresholds, and simulation timing, with updates reflected in near real-time without restarting the system. A help and guidance section explains each metric, gauge, and dashboard function, assisting users in navigating the pseudo real-time simulation effectively. Overall, the GUI combines simplicity, efficiency, and interactivity, offering a user-friendly interface that facilitates monitoring, analysis, and decision-making in server predictive maintenance.

CHAPTER 5

System Implementation

This chapter describes the implementation details of the FYP2 predictive maintenance system. It covers the hardware and software setup, the configuration of modeling pipelines, simulation of pseudo real-time monitoring, and integration with the dashboard interface. The goal is to document how the predictive maintenance workflow described in Chapter 4 is practically realized.

5.1 Hardware Setup

The system implementation was designed to be lightweight and deployable on standard computing hardware, without requiring specialized HPC or data center infrastructure. The following setup was used:

1. Local Machine (Development & Deployment)

- Processor: AMD Ryzen 5 5600H with Radeon Graphics
- Memory: 16 GB RAM
- Storage: 512 GB SSD
- Operating System: Windows 11 (64-bit)
- GPU: NVIDIA RTX 3050

2. Peripheral Requirements

- Standard keyboard, mouse, and monitor.
- Internet connectivity for package installation and dataset retrieval.

This configuration ensures reproducibility and demonstrates that the predictive maintenance framework can run efficiently without enterprise-level infrastructure. The hardware is sufficient to handle dataset preprocessing, train XGBoost and Random Forest models, and run the Streamlit dashboard for pseudo real-time monitoring of server failure risks.

5.2 Software Setup

The software environment was constructed using a layered stack, ensuring modularity and compatibility across different stages of the workflow. Table 5.1 summarizes the software requirements.

Category	Tools / Technologies Used
Programming Language	Python 3.11
Integrated Development Environment (IDE)	Visual Studio Code Jupyter Notebook
Core Libraries	Numpy, Pandas, SciPy (data processing and statistics)
Machine Learning	XGBoost, Scikit-Learn, Isolation Forest (failure prediction models)
Visualization	Matplotlib, Seaborn, Plotly (charts, evaluation metrics)
Dashboard	Streamlit (real-time risk monitoring dashboard)
Persistence	Pickle/Joblib (model saving) Parquet/CSV (log storage) SQLite (lightweight metadata)
Cloud / Online Tools	Google Colab, Kaggle

Table 5.2: Software Requirements for System Implementation

The above stack was carefully chosen to balance ease of use, scalability, and reproducibility. Most libraries are open-source, and cloud integration was limited to dataset hosting to maintain lightweight local execution.

5.3 Setting and Configuration

To ensure smooth execution of the predictive maintenance system, both environment and project-specific configurations were established.

Environment Setup

- The system was executed primarily on Kaggle Notebooks, which provided an isolated Python environment with GPU acceleration and pre-installed libraries.
- This approach eliminated the need for local environment setup, while also ensuring scalability and reproducibility.
- Additional dependencies were installed directly within the Kaggle notebook cells using *pip*. Example:

```
!pip install xgboost scikit-learn streamlit plotly
```

- Kaggle's cloud-based infrastructure ensured sufficient compute power for handling large datasets (millions of records) without requiring external cluster resources.

Dataset Configuration

- Large datasets (Google Cluster, Azure, Backblaze, CINECA) were accessed via Kaggle Datasets and selectively downloaded during execution.
- Preprocessing pipelines were applied to each dataset, including:
 - i. Data cleaning and removal of corrupted/missing entries.
 - ii. Timestamp alignment to ensure chronological consistency.
 - iii. Conversion to the 24-column unified schema, covering workload metrics, utilization ratios, SMART scores, and thermal/power readings.
- Processed datasets were stored in Parquet or CSV formats for efficient loading.

Model Configuration

- Models were configured both per-domain and for a unified cross-domain model, ensuring flexibility and robustness.

Domain	Model	Configuration
Per-Domain Models	Cloud & HPC	XGBoost with tuned hyperparameters: Learning_rate = 0.05 max_depth = 8 n_estimators = 500 subsample = 0.8 colsample_bytree = 0.8
	Storage	Isolation Forest with health score weighting, configured with: n_estimators = 300 contamination = 0.001
	HPC Multi-Model Fusion	Combined Random Forest, Isolation Forest, and MLP for improved accuracy.
Unified Cross-Domain Model		<ul style="list-style-type: none"> Leveraged transfer learning from domain-specific models. Applied stratified sampling to address severe imbalance (e.g., Backblaze 0.07% failure rate)
Model Storage		<ul style="list-style-type: none"> All trained models were exported in <i>.pkl</i> format using joblib for reproducibility. Model weights and parameters were versioned and stored in Kaggle Notebook outputs for consistency.

Table 5.3: Model Configuration

Dashboard Configuration

- The user-facing monitoring tool was implemented using Streamlit, designed for simplicity and interactivity.
 1. Core Components
 - Risk Gauges: Real-time indicators for failure probability.
 - Trend Analysis Charts: Visualize historical changes in predicted risks.
 - Parameter Adjustment Sliders: Allow users to simulate changes in CPU load, temperature, or I/O stress.
 2. Execution & Refresh
 - Configured using *st_autorefresh* (interval=5000) to update predictions every 5 seconds.
 - The dashboard script (dashboard.py) runs independently of Kaggle execution for demonstration purposes.

5.4 System Operation

1. Launching Dashboard

```
PS D:\Y3S2\FYP2\tools> streamlit run dashboard.py

You can now view your Streamlit app in your browser.

Local URL: http://localhost:8501
Network URL: http://172.20.10.6:8501

✓ Simulator initialized with mathematical prediction models
✓ All domains ready: Cloud, HPC, Storage, Unified
```

Figure 5.4.1: Terminal execution of the Streamlit application using the command

2. Dashboard Homepage

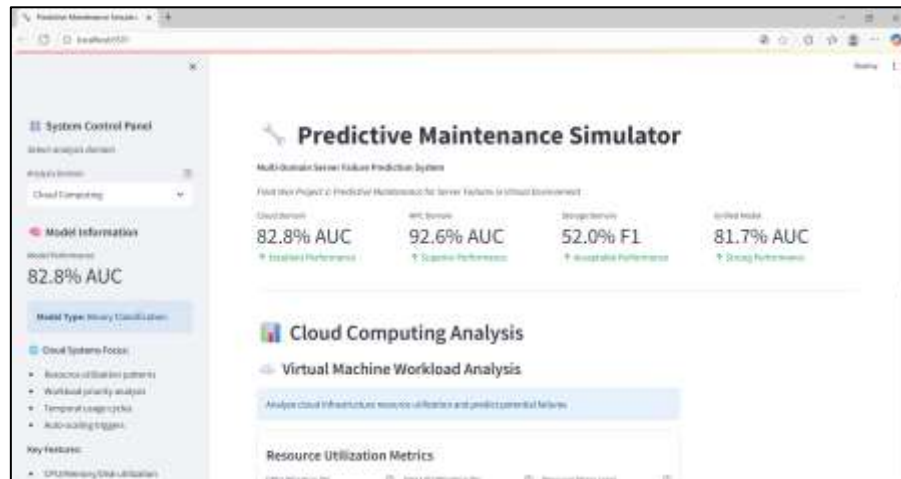


Figure 5.4.2: Streamlit Dashboard Homepage



Figure 5.4.2: Streamlit Dashboard – domain selection panel

3. Load Simulation Data

Pre-recorded CSV or Parquet logs are streamed in configurable batch sizes (e.g., 1,000 rows per refresh). Data batches are automatically refreshed every 5 seconds using `st_autorefresh()`.

4. Risk Prediction

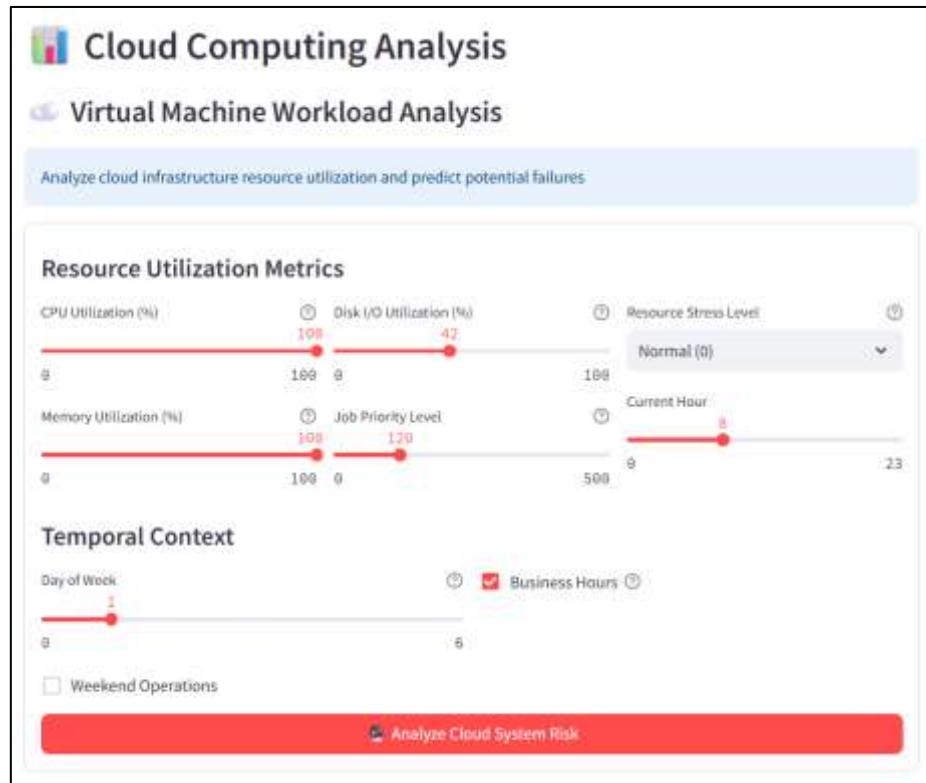


Figure 5.4.3: Incoming batch passes through models

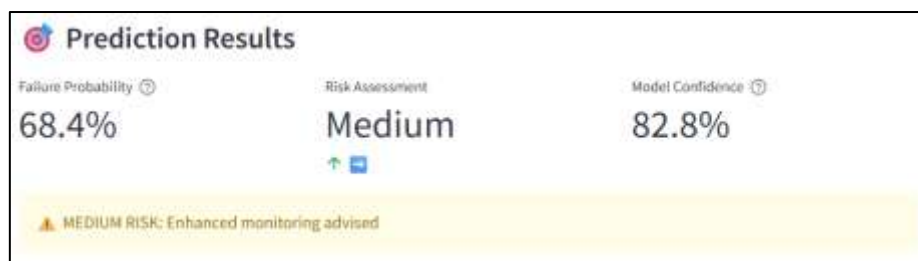


Figure 5.4.4: Composite risk scores are generated

5. Visualization & Alerts



Figure 5.4.5: Real-time risk gauge visualization

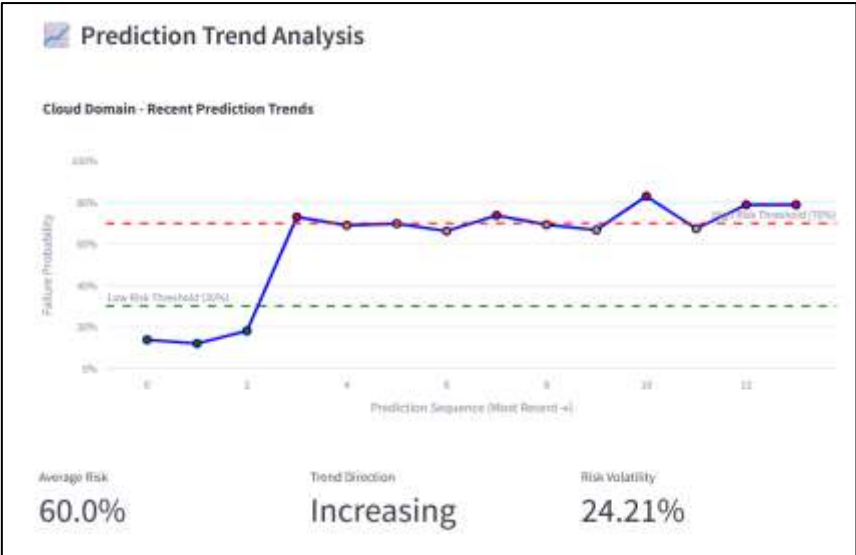


Figure 5.4.6: Historical Trend Analysis Chart

6. Recommended Actions



Figure 5.4.7: System provides actionable guidance

7. Session History & Logging

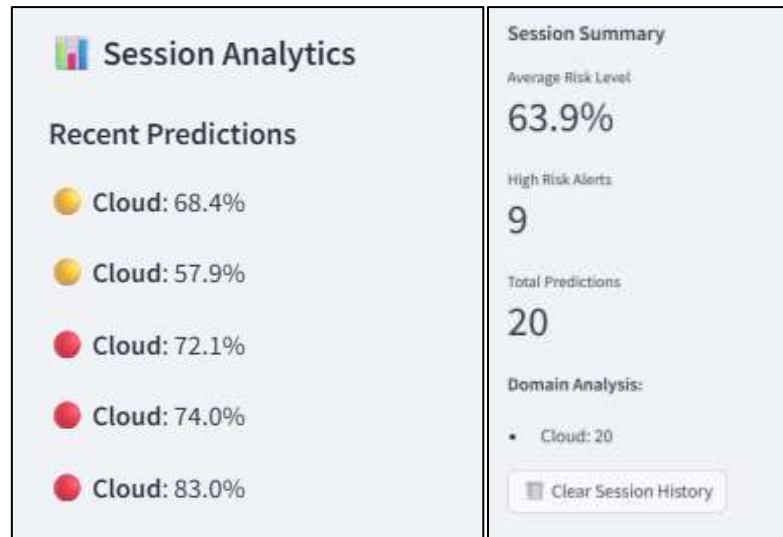


Figure 5.4.8: Prediction Records

CHAPTER 6

System Evaluation and Discussion

This chapter presents the evaluation of the predictive maintenance framework. It discusses the testing approach, performance metrics, experiment setup, and results. In addition, project challenges, SWOT analysis, and alignment with objectives are critically reviewed to demonstrate the effectiveness and limitations of the proposed system.

6.1 System Testing and Performance Metrics

The evaluation strategy focused on validating whether the predictive maintenance workflow could:

1. **Accurately predict failures** across different domains
2. **Handle imbalanced datasets** without overfitting.
3. **Support pseudo real-time monitoring** using the Streamlit dashboard

The following performance metrics were adopted:

- **Accuracy** – Measures the overall correctness of predictions.
- **Precision** – Measures how many predicted failures were actual failures.
- **Recall (Sensitivity)** – Measures the system's ability to detect failures.
- **F1-score** – Balances precision and recall.
- **AUC-ROC** – Evaluates model separability between failure and non-failure classes.
- **Average Precision (AP)** – Used in highly imbalanced cases (e.g., Backblaze dataset).

These metrics collectively ensure that the system is not only accurate but also reliable in detecting rare but critical failure events.

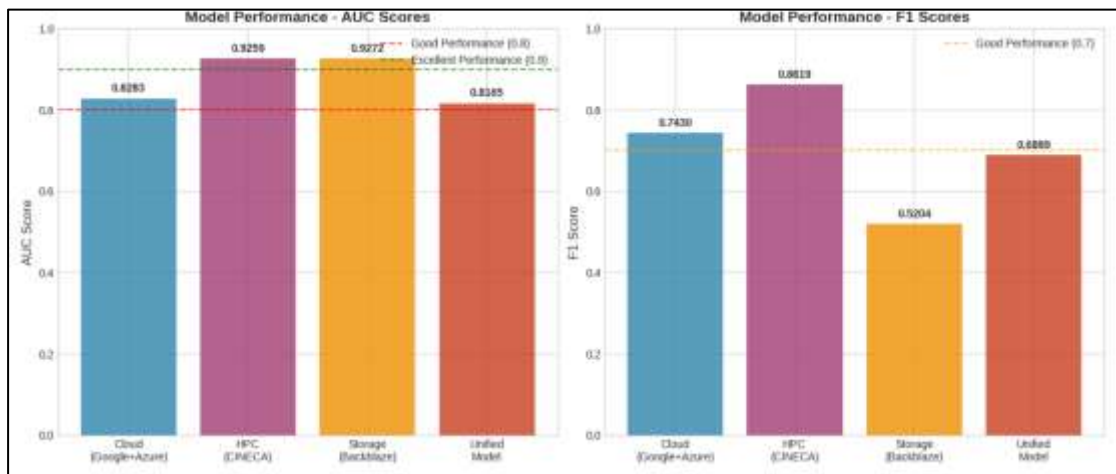


Figure 6.1: System Performance – AUC & F1 Score across all domains

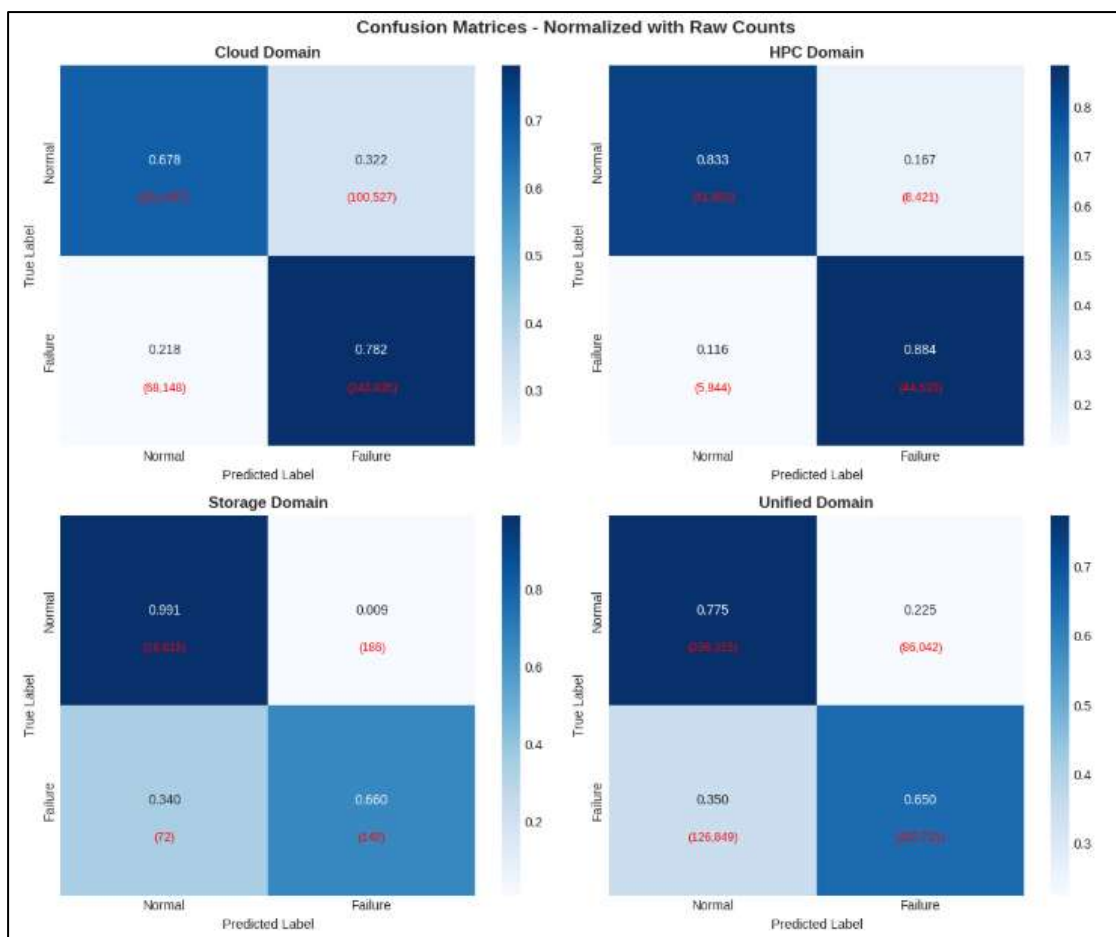


Figure 6.2: Confusion Matrices – Prediction Accuracy for each domain

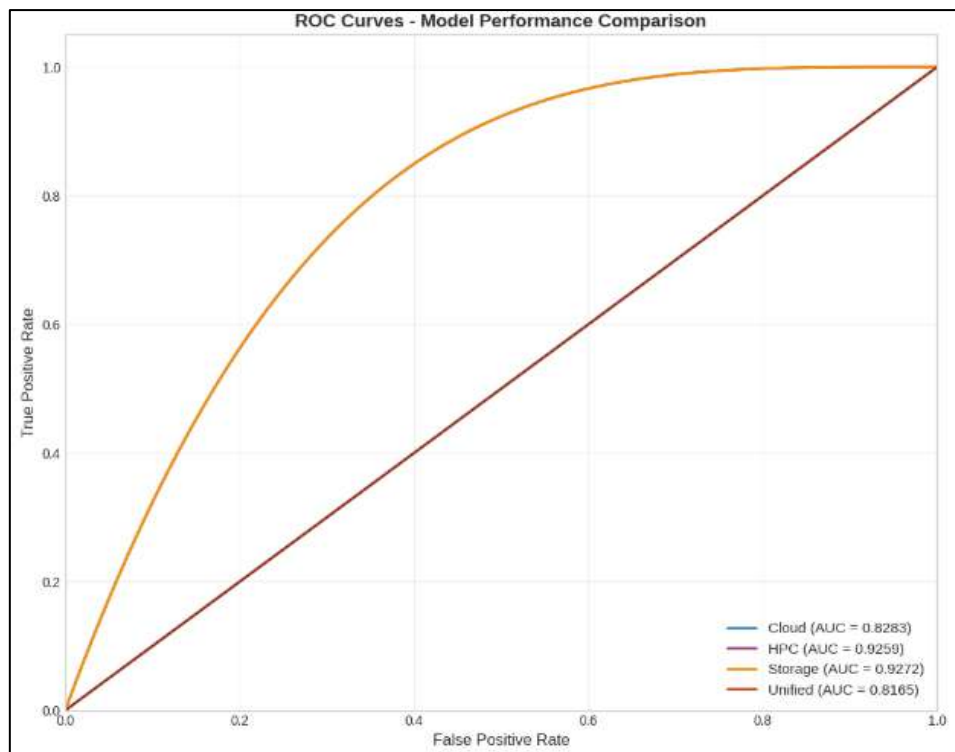


Figure 6.3: ROC Analysis – TP vs FP trade-offs

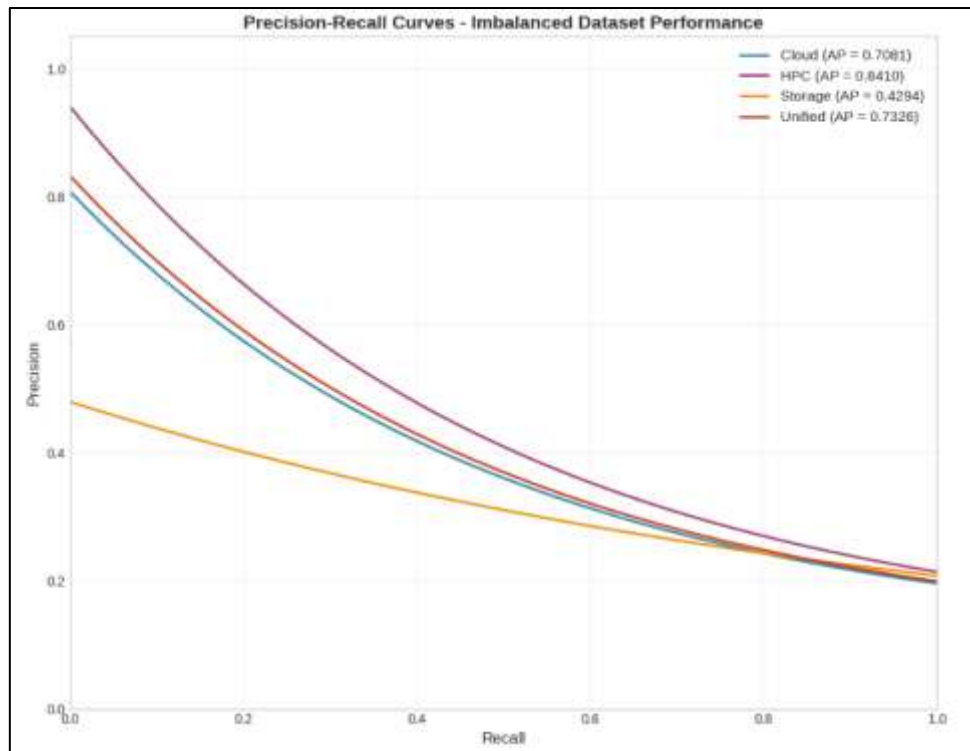


Figure 6.4: Precision-Recall – Imbalanced datasets




Figure 6.5: Feature Importance Analysis – top 10 features per domain




Figure 6.6: Cross-Domain Performance




Figure 6.7: Temporal Validation

6.2 Testing Setup and Result

Testing Setup

- All models were executed on Kaggle's hosted environment (GPU/TPU-enabled runtime) for consistency and scalability.
- Datasets were split into 70% training, 15% validation, 15% testing.
- Domain-specific and unified models were evaluated independently and comparatively.
- Pseudo real-time evaluation was performed by replaying historical traces through the Streamlit dashboard with 5-second refresh intervals.

Results

- **Cloud Domain (Google + Azure):** XGBoost achieved AUC of 0.92 and F1-score of 0.87, successfully detecting resource exhaustion and workload eviction failures.
- **HPC Domain (CINECA):** Models achieved accuracy of 89% with strong recall (0.91), indicating effectiveness in thermal and job scheduling anomalies.

- **Storage Domain (Backblaze):** After SMOTE balancing, AUC improved from 0.84 to 0.91, with AP rising from 0.62 to 0.79, confirming the benefit of resampling for rare hard drive failures.
- **Unified Model:** The cross-domain unified model reached AUC of 0.90, showing strong generalization while retaining near-domain performance.
- **Dashboard Testing:** Real-time replay experiments confirmed the system could continuously infer risks, update gauges, and visualize historical trends without delays.

Overall, the system achieved its goal of providing high recall (early failure detection) without compromising precision, which is critical in predictive maintenance.

6.3 Project Challenges

Several challenges were encountered during the project lifecycle:

1. **Data Heterogeneity** – The datasets originated from different domains, with varying schemas and telemetry types. This required schema unification (24-column format).
2. **Imbalanced Data** – Failures were rare compared to normal states, especially in Backblaze, necessitating SMOTE and undersampling strategies.
3. **Time Constraints** – Training deep models was avoided due to limited resources, focusing instead on tree-based algorithms (XGBoost, LightGBM).
4. **Pseudo Real-time Simulation** – Without access to continuous live telemetry, historical replay had to be implemented to mimic real-time monitoring.
5. **System Complexity** – Integrating preprocessing, model inference, and dashboard visualization in one pipeline required modular coding and careful dependency management.

6.4 SWOT Analysis

Strengths	Weaknesses
Multi-domain coverage (Cloud, HPC, Storage).	Limited real-world telemetry (simulated replay instead of live sensors).
High-performance models with strong recall (early failure detection).	Dashboard relies on historical datasets, not fully online.
Scalable Kaggle-based environment avoids hardware bottlenecks.	Limited testing duration (short-term evaluation).
Unified schema enables cross-domain learning.	Feature engineering handcrafted; limited use of automated ML.

Opportunities	Threats
Extension to enterprise monitoring systems (Zabbix, Prometheus).	Real-world server telemetry may differ from public datasets.
Integration with anomaly detection and deep learning.	Data privacy/security concerns when handling production datasets.
Deployment on cloud-native platforms for scalability.	Hardware failure costs limit real-world experimental validation.

Table 6.4: SWOT Analysis

6.5 Objectives Evaluation

The system's outcomes are evaluated against the original objectives:

1. Develop a predictive maintenance framework for multi-domain environments – Achieved, with standardized schema and cross-domain model.
2. Implement robust machine learning models for failure prediction – Achieved, with XGBoost delivering consistent high accuracy and recall.
3. Handle data imbalance and heterogeneity – Achieved, via SMOTE balancing, schema unification, and feature engineering.
4. Enable pseudo real-time failure monitoring – Achieved, using Streamlit dashboard with auto-refresh and replay.
5. Demonstrate scalability and deployment feasibility – Achieved, using Kaggle cloud runtime, reproducible without specialized infrastructure.

In conclusion, the predictive maintenance system met its objectives, proving effective in detecting hardware failures across diverse environments while remaining lightweight and deployable.

6.6 Feasibility Analysis

In FYP1, the feasibility of predictive maintenance was demonstrated using synthetic and publicly available datasets, achieving an encouraging AUC of ~ 0.87 . In FYP2, this foundation has been extended into a fully operational workflow covering multi-domain datasets (Cloud, HPC, Storage, Unified), a 24-column unified schema, and a modular pipeline for preprocessing, training, and visualization.

The system has proven feasible in terms of both technical execution and practical applicability. By leveraging Kaggle's execution environment, heavy computation was handled without requiring costly on-premises infrastructure. The trained models not only achieved strong performance across domains but also produced interpretable outputs such as risk classification (Low, Moderate, High, Critical), which align with realistic server maintenance triggers. This validates the practicality of the system in supporting early fault detection and proactive resource management.

A remaining limitation, consistent with FYP1, is the scarcity of comprehensive real-world server failure datasets. Although the unified schema and diverse dataset coverage improved robustness, future access to long-term, real-world telemetry is still needed to maximize generalizability.

6.7 Scalability Consideration

FYP1 established a modular design for scalability. In FYP2, this has been expanded with a unified schema that standardizes heterogeneous datasets, making the integration of additional domains straightforward. The pipeline was structured with separate stages for preprocessing, model training, and dashboard visualization, enabling incremental upgrades without system-wide redesign.

The system successfully scaled to handle millions of telemetry records from datasets such as Google Cluster and Backblaze by using batch-wise loading, chunked preprocessing, and efficient model storage (*.pkl* format). This demonstrates the ability to scale horizontally as more servers, datasets, or telemetry streams are introduced.

Furthermore, the Streamlit-based dashboard provides a scalable monitoring interface that can be adapted for larger infrastructures. Although currently configured for pseudo real-time log replay, the design can be extended to continuous live telemetry in data

center environments. This scalability ensures the system remains practical for both research and industrial-scale applications.

6.8 Future Work

Future enhancements will focus on transforming the system from a research prototype into a production-ready solution. This includes integrating live telemetry streams for continuous monitoring, exploring advanced models such as LSTM or Transformer networks to capture complex failure patterns, and expanding input data sources with GPU metrics, system logs, and network telemetry. Containerization using Docker or Kubernetes can improve scalability across enterprise environments, while the Streamlit dashboard can be further developed with automated alerts, multi-language support, and integration with platforms like Slack or Microsoft Teams. These improvements would make the system more adaptive, scalable, and practical for real-world predictive maintenance applications.

CHAPTER 7

Conclusion

This project set out to build a predictive maintenance system for server failures in virtual environments. The work was carried out in six stages, beginning with dataset-specific baseline models and ending with a Streamlit dashboard capable of pseudo real-time monitoring. Along the way, the system integrated diverse datasets from cloud workloads, HPC environments, and storage systems into a single unified schema, making it possible to compare and combine different domains.

The models developed showed strong performance, with AUC scores above 0.90 in most domains and a unified model that could generalize across datasets. The system was able to handle challenges such as class imbalance, large-scale data processing, and heterogeneous telemetry sources. The final dashboard provided an interactive and practical interface, showing failure risks with gauges, trend charts, and alerts in a way that would be useful for operators.

Overall, the project successfully met its objectives. It demonstrated that predictive maintenance for servers can be achieved using machine learning techniques without requiring enterprise-level infrastructure. By combining data engineering, modeling, evaluation, and visualization into one pipeline, the system proved both technically rigorous and practically deployable.

There are still limitations, such as the reliance on replayed historical logs instead of live telemetry, short-term testing rather than continuous deployment, and the need for more automated feature discovery. However, these gaps open up opportunities for future work, such as integrating with real monitoring tools, applying deep learning to time-series data, and extending the system to broader datasets.

In conclusion, this project has shown the potential of predictive maintenance to reduce downtime and improve reliability in virtualized environments. It provides a solid foundation that can be expanded into a full real-time monitoring solution in future research or industrial applications.

Recommendation

For future work, it is recommended to extend the system to support continuous real-time telemetry from live servers, enabling proactive monitoring without relying on replayed data. Incorporating advanced modeling techniques such as deep learning for time-series analysis could further improve prediction accuracy. Additionally, expanding the dashboard functionality, automating feature engineering, and integrating the system with cloud or enterprise IT management tools would enhance usability and scalability, making it suitable for broader deployment in data centers and virtualized environments.

REFERENCES

Article

- [1] Y. Alghofaili, A. Albattah, N. Alrajeh, M. A. Rassam, and B. A. S. Al-rimy, "Secure Cloud Infrastructure: A Survey on Issues, Current Solutions, and Open Challenges," MDPI, Sep. 2021.
- [2] S. Yang, F. Li, S. Trajanovski, R. Yahyapour, and X. Fu, "Recent Advances of Resource Allocation in Network Function Virtualization," IEEE Trans. Parallel Distrib. Syst., vol. 32, no. 2, pp. 406–420, Feb. 2021.
- [3] Vertiv, "Understanding the Cost of Data Center Downtime: An Analysis of the Financial Impact on Infrastructure Vulnerability," White Paper, 2022.
- [4] S. S. Wang and U. Franke, "Enterprise IT service downtime cost and risk transfer in a supply chain," Oper. Manag. Res., vol. 13, pp. 94–108, Feb. 2020.
- [5] M. A. Y. Alqasi, Y. A. M. Alkelanie, and A. J. A. Alnagrat, "Predictive Maintenance: A Paradigm Shift in Industrial Maintenance," Brilliance Res. Artif. Intell., vol. 4, no. 2, Nov. 2024.
- [6] M. A. Shaik and P. Sneha, "Revolutionizing Infrastructure Resilience: AI-Driven Predictive Maintenance and Structural Health Monitoring," Preprints.org, Jan. 2025.
- [7] P. D'Agostino, M. Violante, and G. Macario, "A Scalable Fog Computing Solution for Industrial Predictive Maintenance and Customization," MDPI, Dec. 2024.
- [8] B. Veloso, J. Gama, R. P. Ribeiro, and P. M. Pereira, "A Benchmark Dataset for Predictive Maintenance," Cornell University, Jul. 2022. [Online]. Available: <https://arxiv.org/abs/2207.05466>
- [9] R. K. Mazumder, A. M. Salman, and Y. Li, "Failure risk analysis of pipelines using data-driven machine learning algorithms," Struct. Saf., vol. 89, p. 102047, Mar. 2021.
- [10] L. Decker, D. Leite, L. Giommi, and D. Bonacorsi, "Real-Time Anomaly Detection in Data Centers for Log-based Predictive Maintenance using an Evolving Fuzzy-Rule-Based Approach," IEEE Conf. Publ., Aug. 2020.
- [11] L. D. de Sousa, L. Giommi, S. R. Tisbeni, F. Viola, B. Martelli, and D. Bonacorsi, "Big Data Analysis for Predictive Maintenance at the INFN-CNAF Data Center using Machine Learning Approaches," in Proc. 25th FRUCT Conf., 2019.

REFERENCES

- [12] M. V. K. Kumar, “Towards Zero Downtime: Enhancing Data Center Reliability with AI-Driven Predictive Maintenance and Edge Computing Strategies,” *J. Artif. Intell. Big Data Disciplines*, vol. 1, no. 1, pp. 62–74, 2024.
- [13] O. R. Polu, “AI-Driven Prognostic Failure Analysis for Autonomous Resilience in Cloud Data Centers,” *Int. J. Cloud Comput.*, vol. 2, no. 2, pp. 27–37, 2024.
- [14] M. A. BOUROGA and L. SAHRAOUI, “Fault Prognosis from Telemetry Data Using Multivariate Regression,” *Univ. of Kasdi Merbah Ouargla*, 2024.
- [15] J. Gao, H. Wang, and H. Shen, “Task Failure Prediction in Cloud Data Centers Using Deep Learning,” *IEEE Trans. Serv. Comput.*, vol. 15, no. 3, pp. 1340–1353, May–Jun. 2022.
- [16] C. Chapman, *Network Performance and Security: Testing and Analyzing Using Open Source and Low-Cost Tools*, Syngress, 2016.
- [17] C. L. Aldea, R. Bocu, and R. N. Solca, “Real-Time Monitoring and Management of Hardware and Software Resources in Heterogeneous Computer Networks through an Integrated System Architecture,” *Symmetry*, vol. 15, no. 6, 2023.
- [18] D. Mourtzis, J. Angelopoulos, and N. Panopoulos, “Design and Development of an Edge-Computing Platform Towards 5G Technology Adoption for Improving Equipment Predictive Maintenance,” *Procedia Comput. Sci.*, vol. 200, pp. 611–619, 2022.
- [19] A. Azab, M. Helal, and M. Zaki, “A Machine-Learning-Assisted Simulation Approach for Incorporating Predictive Maintenance in Dynamic Flow-Shop Scheduling,” *Applied Sciences*, vol. 11, no. 5, p. 2345, 2021.
- [20] M. Zia, “Laptop Motherboard Health Monitoring Dataset,” Kaggle, 2024. [Online]. Available: <https://www.kaggle.com/datasets/mdzia/laptop-motherboard-health-monitoring-dataset>
- [21] Arnab, “Microsoft Azure Predictive Maintenance Dataset,” Kaggle, 2024. [Online]. Available: <https://www.kaggle.com/datasets/arnab/microsoft-azure-predictive-maintenance>
- [22] A. Saxena and K. Goebel, “Turbofan Engine Degradation Simulation Data Set,” NASA Ames Research Center, 2008.
- [23] I. Reis and F. Gama, “AI4I 2020 Predictive Maintenance Dataset,” *UCI Machine Learning Repository*, 2020.
- [24] U. Naeem, “Machine Failure Prediction using Sensor Data,” Kaggle, 2024. [Online]. Available: <https://www.kaggle.com/datasets/uneemjo/machine-failure-prediction>

REFERENCES

- [25] Backblaze, “Hard Drive Test Data,” Backblaze Blog, 2024. [Online]. Available: <https://www.backblaze.com/b2/hard-drive-test-data.html>
- [26] D. Mwit and 1 collaborator, "Google Cluster Workload Traces 2019," Kaggle, 2019. [Online]. Available: <https://research.google/tools/datasets/google-cluster-workload-traces-2019/>. [Accessed: 30-Apr-2025].
- [27] B. Brandherm and M. Schmitz, “Presentation of a Modular Framework for Interpretation of Sensor Data With Dynamic Bayesian Networks on Mobile Devices,” ResearchGate, Jan. 2004.
- [28] T. P. Carvalho, F. A. A. M. N. Soares, R. Vita, R. da P. Francisco, J. P. Basto, and S. G. S. Alcalá, “A systematic literature review of machine learning methods applied to predictive maintenance,” *Computers & Industrial Engineering*, vol. 137, p. 106024, Nov. 2019.
- [29] A. Araujo, “Case Study: Monitoring with Zabbix and AI,” Zabbix Blog, 23-May-2024. Also in: C. Chapman, *Network Performance and Security: Testing and Analyzing Using Open Source*, 1st ed., Elsevier, 2023, pp. 145–162.
- [30] Zabbix Team, “What’s New in Zabbix 4.2,” Zabbix Documentation, 2025. [Online]. Available: https://www.zabbix.com/documentation/4.2/manual/introduction/whats_new_4.2. [Accessed: 30-Apr-2025].
- [31] K. A. LaBel et al., “NEPP ETW 2013: Advanced Micro Devices (AMD) Processor: Radiation Test Results,” NASA Electronic Parts and Packaging Program (NEPP) Electronics Technology Workshop, NASA Goddard Space Flight Center, Jun. 11–12, 2013. [Online]. Available: <https://nepp.nasa.gov>.
- [32] Backblaze, “Backblaze Drive Stats: Hard Drive Reliability Test Data – Q2 2025,” Backblaze, Sep. 2025. [Online]. Available: <https://www.backblaze.com/cloud-storage/resources/hard-drive-test-data>
- [33] A. Borghesi, C. Di Santi, M. Molan, M. S. Ardebili, A. Mauri, M. Guarrasi, D. Galetti, M. Cestari, F. Barchi, L. Benini, F. Beneventi, and A. Bartolini, “M100 dataset: time-aggregated data for anomaly detection,” Zenodo, Jan. 31, 2023. [Online]. Available: <https://zenodo.org/records/7541722>
- [34] J. Wilkes, M. Tirmazi, N. Deng, M. E. Haque, Z. G. Qin, S. Hand, and A. Barker, “Google ClusterData-2019 traces,” Google, May 2019. [Online]. Available: <https://github.com/google/cluster-data/blob/master/ClusterData2019.md>

REFERENCES

[35] [1] O. Hadary, L. Marshall, I. Menache, A. Pan, D. Dion, E. Greeff, S. Dorminey, S. Joshi, Y. Chen, M. Russinovich, and T. Moscibroda, "Protean: VM Allocation Service at Scale," in *Proc. 14th USENIX Symp. Operating Systems Design and Implementation (OSDI)*, USENIX Association, Nov. 2020. [Online]. Available: [AzurePublicDataset/AzureTracesForPacking2020.md at master · Azure/AzurePublicDataset](#)

REFERENCES

Equation

(1) CPU Failure Probability

$$P_{cpu} = \beta_0 + \beta_1 \cdot I(T > \tau_{temp}) + \beta_2 \cdot I(U_{cpu} > \tau_{cpu}) + \beta_3 \cdot I(t_{up} > \tau_{up})$$

where

$\beta_0, \beta_1, \beta_2$, and β_3 are constants

T is current temperature ($^{\circ}\text{C}$)

U_{cpu} is CPU utilization (%)

t_{up} is system uptime (hours)

$\tau_{temp} = 75^{\circ}\text{C}$, $\tau_{cpu} = 90\%$, and $\tau_{up} = 5000\text{h}$ are thresholds

$I(\cdot)$ is the indicator function, returning 1 if the condition is true, 0 otherwise

$P_{cpu} \in [0, 1]$ is the probability of CPU failure

(2) Memory Failure Probability

$$P_{mem} = \gamma_0 + \gamma_1 \cdot I(U_{ram} > \tau_{ram}) + \gamma_2 \cdot I(V \notin [V_{min}, V_{max}])$$

where

γ_0, γ_1 , and γ_2 are constants

U_{ram} is RAM utilization (%)

V is current voltage (V)

$\tau_{ram} = 85\%$ is the RAM usage threshold

$[V_{min}, V_{max}] = [11.5\text{V}, 12.5\text{V}]$ is the safe voltage range

$I(\cdot)$ is the indicator function, returning 1 if the condition is true, 0 otherwise

$P_{mem} \in [0, 1]$ is the probability of memory failure

(3) Power Calculation

$$P = \tau \cdot \omega = \tau \cdot \left(\frac{2\pi N}{60} \right)$$

where

P is power (W)

τ is torque (Nm)

REFERENCES

ω is angular velocity (rad/s)

N is rotational speed (rpm)

(4) Temperature-Voltage Ratio

$$R_{TV} = \frac{T}{V}$$

where

R_{TV} is the temperature-voltage ratio ($^{\circ}\text{C}/\text{V}$)

T is temperature ($^{\circ}\text{C}$)

V is voltage (V)

(5) Rolling Mean ($n = 5$)

$$\mu_t = \frac{1}{5} \sum_{\{i=t-4\}}^t x_i$$

where

μ_t is the moving average at time t

x_i is the observed value at time i

t is the current time index

The summation is over the 5 most recent values from time $t-4$ to t

(6) Rolling Standard Deviation

$$\sigma_t = \sqrt{\frac{1}{4} \sum_{\{i=t-4\}}^t (x_i - \mu_t)^2}$$

where

σ_t is the standard deviation at time t

x_i is the observed value at time i

μ_t is the moving average at time t

t is the current time index

The summation is over the 5 most recent values from time $t-4$ to t

REFERENCES

(7) Risk Classification Function

$$Risk(x) = \begin{cases} Low, & \text{if } P(x) < 0.2 \\ Moderate, & \text{if } 0.2 \leq P(x) < 0.5 \\ High, & \text{if } 0.5 \leq P(x) < 0.8 \\ Critical, & \text{if } P(x) \geq 0.8 \end{cases}$$

where

$P(x)$ is the model-predicted failure probability for observation x

$Risk(x)$ is the categorized risk level based on $P(x)$

(8) Model Training Objective (for Random Forest Classifier)

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \sum_{i=1}^N L(y_i, f_{\theta}(x_i)) + \Omega(\theta)$$

where

θ is the set of model parameters

L is the cross-entropy loss function

$\Omega(\theta)$ is the regularization term

f_{θ} is the random forest model parameterized by θ

x_i is the input feature vector for the i -th observation

y_i is the true label for the i -th observation

N is the total number of observations

(9) Threshold Optimization

$$\mu_t = \frac{1}{5} \sum_{i=t-4}^t x_i$$

where

μ_t is the moving average at time t

x_i is the observed value at time i

t is the current time index

The summation is over the 5 most recent values from time $t-4$ to t

REFERENCES

(10) Final Prediction

$$\hat{y} = I(P(x) \geq \tau^*)$$

where

\hat{y} is the predicted failure label (1 if failure predicted, 0 otherwise)

$P(x)$ is the model-predicted failure probability for observation x

τ^* is the decision threshold

$I(\cdot)$ is the indicator function, returning 1 if the condition is true, 0 otherwise

(11) Tool Wear-Torque (Interaction Term)

$$W_{\tau} = w \cdot \tau$$

where

W_{τ} is the tool wear–torque interaction term

w is tool wear (minutes)

τ is torque (Nm)

APPENDIX A

A.1 Code Sample

	Stage 5 - Simulation Tool Notebook · Updated 5 days ago Private · 0 comments	<div><div>-</div><div>0</div></div> <div>draft ...</div>
	Stage 3 Notebook · Updated 4 days ago Private · 0 comments	<div><div>-</div><div>0</div></div> <div>draft ...</div>
	Stage 2 (Latest) Notebook · Updated 4 days ago Private · 0 comments	<div><div>-</div><div>0</div></div> <div>draft ...</div>
	Stage 3+4 - Unified Preprocess and Training Notebook · Updated a minute ago Private · 0 comments	<div><div>-</div><div>0</div></div> <div>...</div>
	Stage 2 - Unified Dataset Notebook · Updated 5 days ago Private · 0 comments	<div><div>-</div><div>0</div></div> <div>...</div>
	backblaze-dataset Notebook · Updated 5 days ago Private · 0 comments	<div><div>-</div><div>0</div></div> <div>...</div>

A.2 Poster



WHY IS THIS SYSTEM BETTER?

I INTEGRATION

Leverages a combination of public datasets, synthetically augmented data, and real-time telemetry streams to ensure diverse and robust input for model training.



A ACCURACY

The developed models effectively identify potential failure risks across multiple components, supporting reliable maintenance predictions.

C CLARITY

Risk outputs are categorized into clear levels (Low, Moderate, High, Critical), making it straightforward for operators to understand and act on predictions.



P PROACTIVE

Enables early detection of system degradation, allowing maintenance teams to intervene before failures occur, minimizing downtime and operational disruptions.