PKM-DOC-21-0007
Revision 0

Integrated Risk-Informed Condition Based Maintenance Capability and Automated Platform: Technical Report 3

Advanced Reactor Development Projects

7/15/2021

Matthew Yarlett, Brad Diggans, Nicholas Goss
PKMJ Technical Services, LLC

Vivek Agarwal, Andrei V. Gribok, Vaibhav Yadav, James A. Smith, Nancy Lybeck, and Koushik A. Manjunatha
Idaho National Laboratory

Harry Palas
Public Service Enterprise Group (PSEG) Nuclear, LLC

PKMJ Technical Services, LLC
1000 Westinghouse Drive, Suite 128
Cranberry Township, PA, 16066

Page intentionally left blank

ABSTRACT

This project is a collaborative research effort between PKMJ Technical Services LLC, Idaho National Laboratory, and Public Service Enterprise Group (PSEG) Nuclear, LLC. The collaboration, led by PKMJ Technical Services LLC, is part of the industry Funding Opportunity Announcement (FOA) award under Advanced Nuclear Technology Development FOA #DE-FOA-0001817. The pilot demonstration focuses on the Circulating Water System (CWS), an important non-safety-related system that impacts the power generation capability of the plant site. Achieving risk-informed condition-based Predictive Maintenance (PdM) on the CWS will result in significant economic benefits, and the developed methodologies can also be applied to other plant systems. This approach supports an industry goal of ensuring that nuclear power generation remains a viable, economically competitive option in the energy market. Operation and Maintenance (O&M) costs include labor-intensive Preventive Maintenance (PM) programs that involve manually performed inspection, calibration, testing, and maintenance of plant assets at periodic frequencies as well as time-based replacement of assets, irrespective of condition. This project offers an alternative by focusing on risk-informed condition-based maintenance to reduce O&M costs while still maintaining plant health and safety.

This report summarizes the progress made toward achieving a risk-informed condition-based maintenance approach. The research and development (R&D) activities presented in this report are associated with development of a nuclear digital platform application, integration of fault signature models, and automated work management processes. The fault signatures and Machine Learning (ML) models are key components in predictive analytics and are heavily leveraged to improve the insights received by existing plant process data sources. Availability of the analysis results within a centralized digital platform enhances efficiency by enabling automation of activities otherwise performed manually. Personnel are presented with enhanced information that can be used to evaluate plant status and risks. Utilizing the enhancements to data analytics supports automated responses, (i.e. issuance of work orders) to address developing equipment faults and thus preventing forced, unplanned shutdowns of components or systems.

The R&D activities described within this report lay the foundation for developing and demonstrating a digital automated platform to centralize the implementation of condition monitoring and response to equipment faults. The digital automated platform is cloud-based and designed to enable improved efficiency of plant processes. The digital platform includes content related to maintenance optimization, fault signature analysis, and plant records, which can all be used to support efficiencies when located within a centralized digital platform. These efficiencies could be further enhanced when deployed through industry-wide deployment of the technology to improve insights and processes based upon economies of scale.

Page intentionally left blank

EXECUTIVE SUMMARY

In support of the U.S Department of Energy (DOE) Office of Nuclear Energy's (NE's) priority for Advanced Nuclear Technology Development, PKMJ Technical Services LLC, in partnership with Idaho National Laboratory and Public Service Enterprise Group (PSEG) Nuclear, LLC is leading this research to address challenges in the implementation of risk-informed, condition-based Predictive Maintenance (PdM). The research outcomes will consist of models and methods to enable deployment of a risk-informed PdM program at a Nuclear Power Plant (NPP). Implementation of a risk-informed PdM program is critical for ensuring long-term safe and economical operation, automation, efficiency, and enhanced reliability of plant systems in NPPs.

To achieve the project objective, three goals are defined. They are:

Goal 1: Develop a risk-informed approach to optimize equipment maintenance frequency: Perform R&D activities to develop a new capability that enables the optimization of Preventive Maintenance (PM) frequencies for the Circulating Water System (CWS) components based on a risk-informed approach. In this activity, historical plant process data for the CWS, PM and Corrective Maintenance (CM) records, failure data, and preexisting expert opinions will be utilized to enhance the risk insights needed to prioritize and inform maintenance decision making.

Goal 2: Develop a risk-informed condition-based maintenance approach: Perform R&D activities using advancements in sensor technologies and advanced data analytics to develop and deploy digital monitoring and to develop an automated diagnosis and prognosis process that offers insight into the health of components within the CWS.

Goal 3: Develop and demonstrate a digital, automated platform to centralize the implementation of monitoring technologies: Perform R&D activities to integrate the capabilities developed in Goals 1 and 2 into a centralized automated platform to support the broadest possible implementation of technologies for use by industry to achieve the greatest returns on investment based on economies of scale.

The outcomes presented in this report directly address Goal 3 and describe improvements and progress made since the completion of Goals 1 and 2. These include:

1. Implementation of the PKMJ digital platform based on requirements that align with nuclear industry requirements. This will provide guidance to NPPs, as they plan to develop or acquire their own digital platforms;
2. Development of intuitive user interface to review information in a single-page application and make decisions related to maintenance strategy optimization and component faults;
3. Approaches to data extraction, cleansing, and enhancement, which are necessary for transforming the raw data received in order to properly stage the data for use in CBM modeling;
4. Enhancement of CBM models to better identify component faults;
5. Approach, assumptions, execution, and results of risk-informed models based on the three-state Markov chain process;
6. Development of Advanced Pattern Recognition (APR) models for comparison with CBM models; and

7. Design and implementation of an automated work management process that enables utilities to automatically respond to component faults by planning corrective maintenance.

The report highlights the cloud-hosted digital platform developed for this project as a centralized location in which utility personnel can perform tasks such as maintenance strategy optimization and the evaluation of component faults. The insights identified within the digital platform are valuable for NPPs as they implement new technologies to reduce costs and improve processes. The PKMJ Nuclear Digital Platform utilizes cloud services, Artificial Intelligence (AI) / Machine Learning (ML), and Natural Language Processing (NLP), etc., to provide unique, effective services to the nuclear industry.

Page intentionally left blank

ACKNOWLEDGMENTS

This report was made possible through funding by the United States Department of Energy's Opportunities for Advanced Nuclear Technology Development FOA #DE-FOA-0001817. We are grateful to Alison Hahn of the United States Department of Energy as well as Bruce P. Hallbert and Craig A. Primer at Idaho National Laboratory for championing this effort. We thank John M. Shaver and Nikki Peterson at Idaho National Laboratory for technical editing of this report. We would also like to thank the following staff members at PKMJ Technical Services LLC, PSEG Nuclear, INL, and GrayMatter for their valuable technical contributions:

- PKMJ Technical Services, LLC
    - Blaze Sterling
    - Brad Anthony
    - Casey Emery
    - Eric Knop
    - Frank Koski
    - Kiranmayi Kale
    - Mathew Mackay
    - Matthew Broman
    - So Kim
    - Stephen Dolansky
    - Steven Yurkovich
- Public Service Enterprise Group Nuclear, LLC
    - Matthew Pennington
    - James Morocco
- INL
    - Michael A. Ziccarelli
    - Cody M. Walker
- GrayMatter
    - Andrew Drake
    - Mitchell Keller
    - Philip Henri

Page intentionally left blank

CONTENTS

FIGURES

TABLES

Page intentionally left blank

ACRONYMS

ADLS	Azure Data Lake Storage
AI	Artificial Intelligence
API	Application Programming Interface
APR	Advanced Pattern Recognition
ARIMA	Auto Regressive Integrated Moving Average
AUC	Area Under Curve
AWS	Amazon Web Services
BOP	Balance of Plant
CBM	Condition-Based Maintenance
CDF	Cumulative Density Function
CI/CD	Continuous Integration / Continuous Deployment
CRADA	Cooperative Research and Development Agreement
CM	Corrective Maintenance
CW	Circulating Water
CWP	Circulating Water Pump
CWS	Circulating Water System
DOE	Department of Energy
DT	Differential Temperature
EPRI	Electric Power Research Institute
EMS	Enterprise Management System
FaaS	Functions as a Service
FEG	Functional Equipment Group
FN	False Negative
FOA	Funding Opportunity Announcement
FP	False Positive
FME	Foreign Material Exclusion
GUI	Graphical User Interface
INL	Idaho National Laboratory
JWT	JSON Web Token
INL	Idaho National Laboratory
LF	Line Frequency
LWR	Light-Water Reactor
LWRS	Light-Water Reactor Sustainability
M&P	Motor and Pump

MIB	Motor Inboard Bearing
ML	Machine Learning
MME	Measuring and Monitoring Equipment
MOB	Motor Outboard Bearing
MP	Maintenance Plan
NDP	Nuclear Digital Platform
NLP	Natural Language Processing
NPP	Nuclear Power Plant
O&M	Operation and Maintenance
OID	Object ID
PDF	Probability Density Function
PdM	Predictive Maintenance
PKMJ	PKMJ Technical Services, LLC
PM	Preventive Maintenance
PMO	Preventive Maintenance Optimization
PRA	Probabilistic Risk Assessment
PSEG	Public Service Enterprise Group
RBAC	Role-Based Access Control
R&D	Research and Development
ROC	Receiver Operating Characteristic
RV	Random Variable
SaaS	Software as a Service
SAS	Shared Access Signature
SHAP	Shapley Additive Explanations
SME	Subject Matter Experts
SSL	Secure Sockets Layer
TERMS	Technology Enabled Risk-informed Maintenance Strategy
TN	True Negative
TP	True Positive
VA	Valve Axial
VPF	Vane Pass Frequency
VR	Valve Radial
WO	Work Order
XGBoost	eXtreme Gradient Boosting

Page intentionally left blank

Integrated Risk-Informed Condition-Based Maintenance Capability and Automated Platform

1. INTRODUCTION AND BACKGROUND

The primary objective of this research is to address challenges in the implementation of risk-informed, condition-based Predictive Maintenance (PdM) for identified plant assets in order to reduce operating costs while still maintaining the safety and reliability of commercial Nuclear Power Plants (NPPs). To achieve this objective, risk models integrated with predictive models are being developed by utilizing advancements in data analytics, deep learning, Machine Learning (ML), and Artificial Intelligence (AI). The project is also researching and developing an automated platform to support agile business processes through implementation of technology for use by the nuclear industry. The project outcomes will consist of models and methods that will enable industry-led innovation and technological deployment in the current fleet of U.S. NPPs, ensuring that the nuclear industry remains an economically competitive and viable option in the energy market.

This project is a collaborative research effort between PKMJ Technical Services LLC (PKMJ), Idaho National Laboratory (INL), and Public Service Enterprise Group (PSEG) Nuclear, LLC. This collaboration, led by PKMJ, is part of the industry Funding Opportunity Announcement (FOA) award under Advanced Nuclear Technology Development FOA #DE-FOA-0001817. A pilot demonstration of developed models and methods will be undertaken by PSEG Nuclear, LLC at their Salem Nuclear Power Plant. The pilot demonstration will focus on the Circulating Water System (CWS), an important non-safety-related system that impacts the plant site power generation capability. Achieving risk-informed condition-based PdM for the Salem CWS demonstrates the value of this technology, the potential cost savings, and lays the foundation for expanding the technology to other plant assets.

To achieve the project objective, three goals are defined, as listed below.

Goal 1: Develop a risk-informed approach to optimize equipment maintenance frequencies.

Research and Development (R&D) activities for developing a new capability that enables optimization of Preventive Maintenance (PM) frequencies for the CWS, based on a risk-informed approach. In this activity, information extracted from historical plant process data, PM and Corrective Maintenance (CM) records, failure data, and expert opinions, specific to the CWS, will be utilized to enhance risk insights in order to prioritize and inform maintenance decision making.

Goal 2: Develop a risk-informed condition-based maintenance approach.

R&D activities will be performed using advancements in sensor technologies and advanced data analytics to develop and deploy digital monitoring and develop automated diagnosis/prognosis regarding the health condition of the plant CWS. Using the capability developed through Goal 1, these R&D activities will employ advanced monitoring and diagnostic/prognostic models to recommend the performance of CBM activities on plant equipment. This will move maintenance away from scheduled, frequency-based activities to activities performed only when necessitated by conditions to reduce the amount and types of maintenance performed. This marks the transition to technology-enabled, condition-based, risk-informed maintenance activities.

Goal 3: Develop and demonstrate a digital, automated platform to centralize the implementation of technology monitoring.

The move from time-based maintenance to CBM represents a significant shift in both the methods and tools for plant monitoring and cost reduction. The greatest economies of scale are realized when these technologies are centralized—that is, deployed in multiple plant settings or in a fleet of plants—for monitoring a fleet (or fleets) of components. R&D activities will be performed to integrate the capabilities developed in Goals 1 and 2 into a centralized automated platform to support the broadest possible implementation of technologies (for use by industry) in order to achieve the greatest returns on investment and economies of scale. The platform will automate business processes such as automatically generating work orders (WO's), managing inventory parts, aligning work with the right skilled/trained field workers, and updating the system with feedback received once the work package is complete. The platform will provide a schedule optimization tool to track and realign (if required) activities to ensure on-time completion. This is achieved through the development of applications that interface with one another on the platform while pulling the required utility information from the central data-lake.

Outcomes of the R&D activities including enhancements to the work performed as part of Goal 1 and 2 as well as completion of work through Goal 3 are detailed in this report. The notable outcomes presented in the report include:

- Development of an Azure-based digital platform that allows industry stakeholders to utilize services developed during the research activities. Implementation of the digital platform allows personnel to take advantage of previously unachievable efficiencies such as digital review of services and outputs, improved access to information, and visualizations. These features implemented in the digital platform support the goal of improving work processes and decision making at NPP's.

- Enhancement of ML models using heterogeneous plant process and vibration data, collected from the Salem NPP's CWS. These models are designed to diagnose CWS pump or motor degradation risks based on fault signatures developed using historical process and WO data. Degradation risks are events associated with components (or systems) that may result in the component (or system) operating outside its design specifications (ratings). If a degradation risk is not addressed, it could lead to a component failure, plant derate, or plant trip depending on the function of the degraded component (or system). The developed diagnostic models are a template that can be expanded to other faults and components associated with the CWS. The modeling technique provides the foundation for condition-based monitoring and replacement of plant assets throughout the plant. The replacements of plant assets are currently performed at set time intervals irrespective of the state of the health of the asset which can be enhanced with a condition-based maintenance strategy.

- Formulation of a three-state Markov chain risk model provides probabilities of components being in either an operational or maintenance state. The parameters associated with the rate of transition between different states of the Markov chain models were estimated using WO data for both Salem units. The economic model formulation utilizes the probability values to presents an initial cost-benefit analysis for PdM strategies. The economic analysis includes both time-independent and time-dependent parameter variation, fostering risk-informed decision-making. This provides technical and business justifications to achieve the reduction in maintenance costs by transitioning to a risk-informed CBM strategy.

- Development of the Automated Work Management process, an integral part of the digital platform, enables automation of WO creation. The created WOs detail the activities required to correct degradation risks identified via CBM models. The Automated Work Management process is informed by outcomes of the CBM models and can utilize a utility's Enterprise Management System to create a WO. The report generated by the NDP can also be attached as a reference to documentation generated by the utility. This ability to streamline the process from an event (degradation risk) identification to corrective maintenance (informed by the condition of the plant asset) supports resource (labor and material) optimization, cost savings, and improved response times for plant processes.

1.1 Motivation/Background

Over the years, the nuclear fleet has relied on labor-intensive, time-consuming maintenance programs, driving up operation and maintenance (O&M) costs in order to achieve a high-capacity factor. A well-constructed PdM approach would allow commercial NPPs to reliably transition from current labor-intensive PM programs to a technology-driven PdM program, as shown in Figure 1, thus eliminating unnecessary O&M costs and ensuring the economic competitiveness of the nuclear industry in today's energy market. Right now, these O&M costs are a major contributor to the total operating costs. They involve manually performed inspection, calibration, testing, and maintenance of plant assets at periodic frequencies, along with scheduled replacement of assets, irrespective of their condition. This has resulted in a costly, labor-centric business model. Transition to the technology-centric business model will significantly reduce PM activities and drive down costs, since labor is a rising cost and technology is a declining cost. This transition will also enable NPPs to maintain high capacity factors—perhaps even raise them while still significantly reducing O&M costs.

Figure 1 - Transition from a PM Program to a Risk-Informed PdM Program

The challenges facing the industry are clearly understood by regulators, operators, and vendors alike. The Nuclear Energy Institute has issued several efficiency bulletins related to reducing maintenance costs. The PdM R&D plan [1] laid the foundation for real-time condition assessment of plant assets. Successful execution of the R&D plan will result in the development of a deployable, risk-informed PdM maintenance program for plant use, thus enhancing the safety, reliability, and economics of plant operation.

1.2 Report Layout

This report is organized as follows:

- Section 2 presents architecture and design considerations for the Nuclear Digital Platform (NDP).
- Section 3 describes intended user interfaces within the NDP.
- Section 4 describes the continued research effort associated with condition-based monitoring following Goal 2.
- Section 5 describes the development of an Advanced Pattern Recognition (APR) model to evaluate equipment reliability at Salem NPP.
- Section 6 presents the Automated Work Management Process developed within the NDP.
- Section 7 summarizes the work performed under this project and demonstrates the achievement of the project goals.

2. DIGITAL PLATFORM ARCHITECTURE AND DESIGN

This section describes requirements that were established through development and design of the Nuclear Digital Platform by PKMJ (referred to as PKMJ NDP or NDP throughout this report). The PKMJ NDP was designed to be a centralized solution to support large-volume storage and access, large-volume data processing, advanced data analytics techniques (e.g., AI/ML), data and information visualization, and reporting. PKMJ utilizes enhanced business intelligence techniques in support of NPP customers; however, these tools are used in industries throughout the world. The PKMJ NDP takes input from nuclear industry subject matter experts and combines it with input from mixed discipline teams of data experts in order to apply cutting-edge principles to rapidly explore data, unlock and test new ideas, and turn those ideas into services.

As discussed in the Goal 1 Technical Report [2], PKMJ opted to develop the NDP using Microsoft's Azure Cloud Services. While any Cloud Service Provider can be chosen to support a digital platform, this report covers PKMJ's experience in developing the NDP within the Azure Cloud Environment beyond the details discussed in the first Technical Report [2]. Selection of a cloud environment provides the structure, scalability, security, etc., necessary to deliver enhanced digital service offerings to the industry.

2.1 General Design Concepts

The PKMJ NDP is an Azure Cloud-Based Web Application. The design of the NDP follows recommendations from Microsoft and can be generally divided into the following major components: NDP Web Application, Azure Active Directory, Azure Keystore, Application Programming Interface (API) Management, Azure Functions, Blob Storage, and Databases. These different components are utilized together to provide the functionality required for the NDP. A flow chart of the technology stack used for the NDP is shown in Figure 2 below.

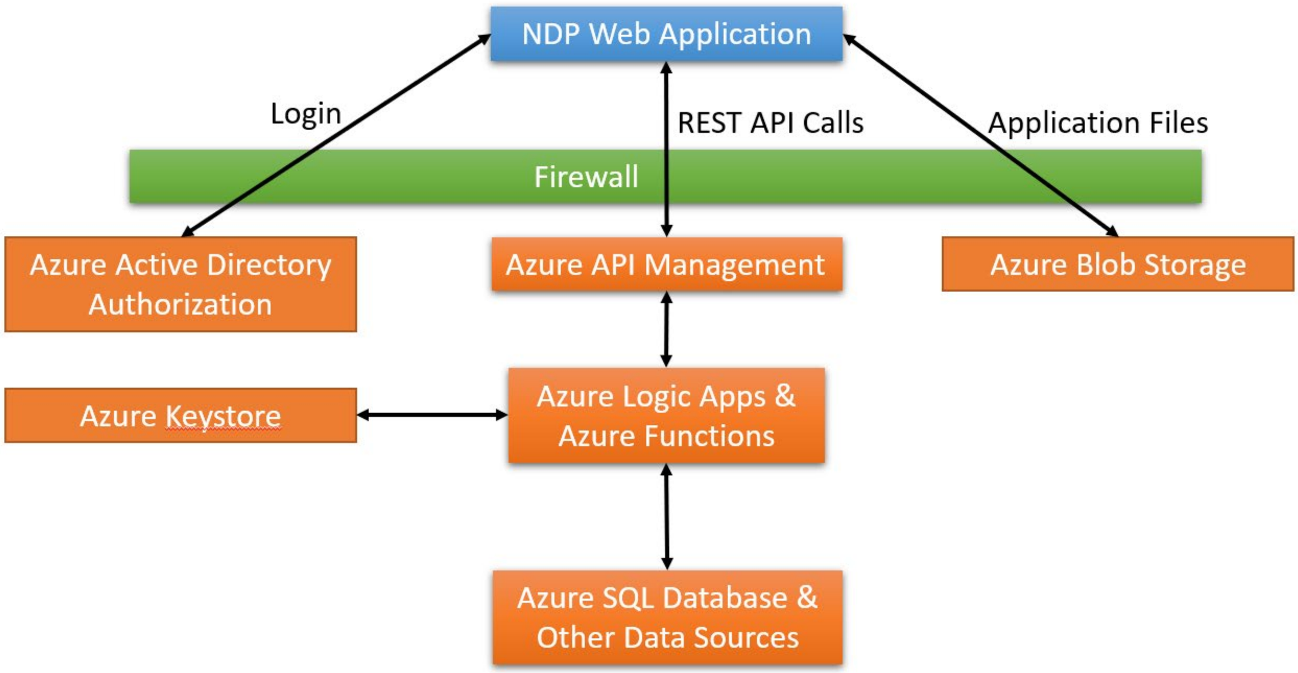

Figure 2 - Web Application Technology Stack & Architecture

The NDP Web Application consists of the React files that render a working application. The files for the application are stored and pulled from Blob Storage to serve up the application. Blob Storage services are also tied to versioning control scheme to automatically receive updates to the web application. When the NDP's front-end makes data requests, Azure API Management provides an intermediary layer that controls data requests, adds granularity, and enhances security. Web application requests are routed to the APIs, which are then routed to the back-end services within the cloud platform. Azure Functions form the back-end services that provide for the application's retrieval and business logic. Azure Keystore provides access credentials for back-end services without needing to hard-code the credentials. The Azure Functions retrieve data from a SQL Database that stores the datasets used for the web application.

The NDP is designed to use cloud resources to provide a high degree of availability and reliability, enabling the platform to perform its function when called upon to do so. Cloud service providers often guarantee high availability for applications built on their respective platforms. For example, Microsoft Azure guarantees a 99.95% availability rate based on its service agreement [3]. This can be further enhanced by using Cross-Regional Redundancy to mitigate against failure of the host data center(s). As currently designed, the NDP does not require immediate response by site personnel to alerts (or degradation risks), so this redundancy was not implemented; however, it is a consideration for the platform as tasks requiring urgent action and review are incorporated into the platform.

The web application is designed to be accessed via various common web browsers such as Google Chrome, Microsoft Edge, or Mozilla Firefox. The NDP functions using Secure Sockets Layer (SSL) protocol meaning that the platform is accessed using an encrypted HTTPS URL. This approach was selected as a secure way to allow access for NDP target end users such as engineers, planners, and monitoring teams. These various work groups typically have convenient access to computers as part of their normal job functions. A future consideration for an industry-wide digital platform may be to support implementation using a tablet to support mobile access within the plant; however, this was not needed in order to perform the plant tasks required of the application in this project. Although a specific mobile application was not required, the standards used for the common web browsers typically support browsers on tablets and other mobile devices.

A central login is provided for all users to minimize the need to log-in multiple times to access the various services integrated into the NDP. This approach is enhanced by using Azure's Active Directory service, which would allow the NDP to authenticate using the end user's single sign-on platform, enabling a secure log-in to the platform. When access is required across organizations, Azure's Active Directory Business to Business collaboration feature should be considered to allow secure guest access to the web application while still maintaining control over corporate data.

The utilization of Azure Active Directory also provides specific controls associated with permissions within the application. End users can be allowed access to security groups based on Active Directory rules (or manual entry) that automatically assigns the user to the appropriate access classification to services within the NDP. This stage of authorization assigns users to their specific fleet and/or site and controls access to data and services available to the fleet and/or site user. Once the user is assigned to the appropriate high-level permission group, user-level permissions are used to control functions within the NDP services. Within a selected service, the permissions or roles will assign the user as Administrator, Approver, Reviewer, or Read Only. Users can, as needed, have different permissions on a per-service basis, depending on their roles or responsibilities.

2.2 Networking and Security/Monitoring

PKMJ's NDP utilizes NPP data to make decisions and support the overall goal of improving plant efficiencies and reducing maintenance costs. As a result of the data included, it is critical that the NDP is secured from malicious actors who may attempt to access and misuse the data. The security design of the NDP follows industry standards and best practices for web application hosting, as discussed in the first Technical Report for this project [2]. These practices include the generation of log files detailing events occurring within the application as well as application of minimal required access security rules to limit access to the minimum level necessary to perform a given function. However, certain specific requirements impacted the NDP development path, as will be examined in detail within this report.

The most impactful security requirement associated with the application was controlling the access that virtual network resources had within the NDP. Early iterations of the virtual network design required the use of external Software as a Service (SaaS) tools, which were thought to have required public access to the internet when utilizing the services. The network team quickly identified that this approach was unacceptable from a security standpoint as data would be available over the internet. It was identified that the resources needed to be configured using Azure Private Endpoints, which allow virtual network resources to privately connect to other resources as if they were part of the same network. This design allows for the use of Azure's backbone to support communication instead of using the internet.

Designing the virtual network resources to use private endpoints challenged the design associated with the SaaS tools intended to be used for the application. Additional work was required to exclude those tools that could not be used securely and to reconfigure those that did support private endpoints. A best practice for teams developing a digital platform is to consider and understand where all traffic is routed to ensure adherence to network security and Information Technology standards for the data within the application. This aspect of data security is critical for preventing intrusions or malicious acts being performed on the data.

Data security was enhanced by ensuring that additional encryption standards were applied within the application architecture. Without strict encryption standards applied, service providers within the cloud-hosted web application may end up with access to the underlying data. A specific example of this is the Azure Databricks service, which was configured to ensure that no public Internet Protocol (IP) addresses could access the data, and that the metastore utilized by Databricks was private. The implementation of additional encryption is a best practice for preventing malicious and/or unauthorized access to the application's underlying data.

Another layer of application security is to ensure that requests made to the application are valid. The NDP utilizes API Management services in Azure to validate JSON Web Token (JWT) files. This layer of validation ensures that functions calling data within a restricted service are restricted based on the user's privilege and the request made. In addition, Azure Active Directory Object ID's (OID's) provide unique identification for application objects in Azure. The OID's serve as another layer of validation to ensure that requests are valid. Where possible, caching can be considered for use in storing recent user information to reduce the overall number of requests. The design of any digital platform must consider how to restrict users from viewing data outside their privilege.

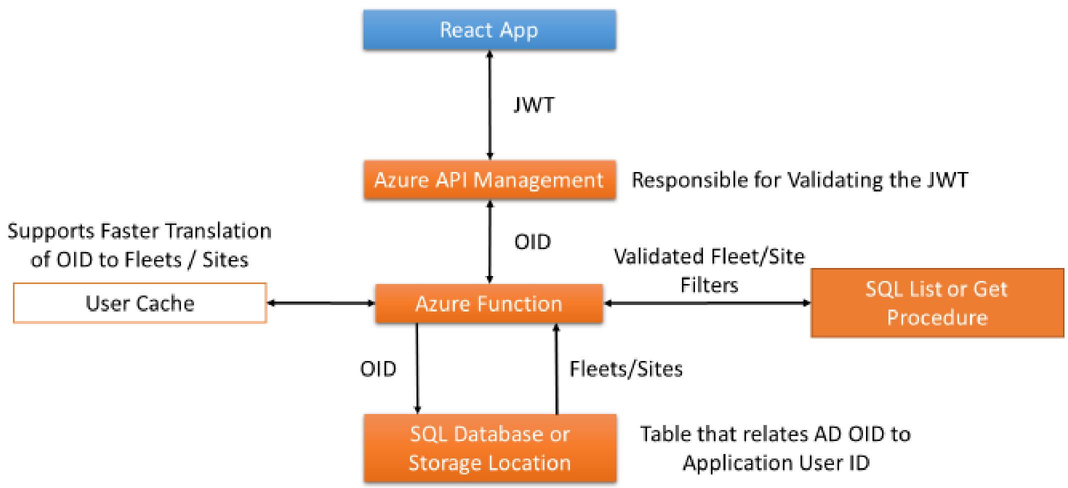

Figure 3 - Data Authorization Structure

The NDP utilizes several datasets to provide users with the right data to extract information and make decisions within the application. These datasets are often obtained from sources such as third-party vendors or the plant itself for plant process data and WO data. Vibration data from KCF Technologies as well as PSEG's PI Historian computer data are utilized by the NDP to evaluate component health. A method utilized by Azure Storage Accounts for allowing access is that of using Shared Access Signature (SAS) API's to grant restricted access rights to Azure Storage Account resources such as blob storage, files, queues, and tables.

The generated SAS token restricts access by service type, resource type, permissions (CRUD), time frame, and IP whitelists [4]. This method of allowing access is problematic as the SAS tokens cannot be fully audited [5]. A review was performed to determine whether SAS keys can be tied to Role Based Access Control (RBAC) security protocols in order to resolve this issue, but the only restriction that worked well involved aligning firewall rules to restrict access to a given IP range. Unfortunately, the IP range for authorization is visible in the key itself, thus defeating this security measure.

Our recommended best practice for resolving the concern with SAS keys is to develop a new API Management protocol architecture that is consistent with the required security protocols [6]. This is done by combining the API Management service with a new, customer Azure Function App to provide the required security. This Azure Function App is generalized for use with any Azure Storage Account to support the various functionalities built into the NDP. This approach also supports scalability for future data inputs to the platform as the redeployment of the Function App is trivial. Figure 4 shows additional detail on the configuration used for the NDP.

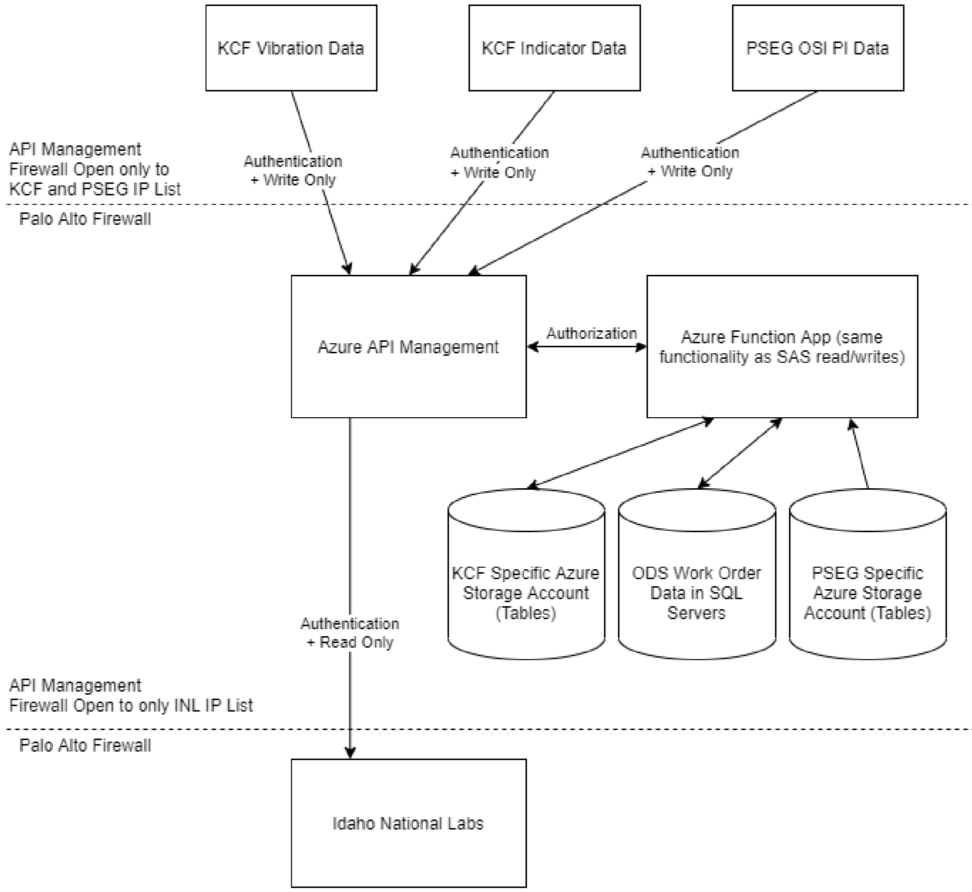

Figure 4 - Custom API Structure for External Data Sources

To meet security standards, security for any digital platform must consider the least required access necessary to perform the desired functions. The developer of any digital platform must consider all interfaces associated with the application (both internal and external) to determine the appropriate security stance. All communication streams should be well understood by the design team to ensure that full control of the application and its data are always maintained.

2.3 Data Structures

Data creates the underlying foundation for any work performed in the NDP as the data drives the analytics and generation of insights, review of insights, as well as resulting actions or services. As the data utilized for the application are diverse and sourced from many different locations, the structures used to store, retrieve, and modify the data must be suitable for the intended application. For example, platform developers must ensure that application data, plant enterprise system data, and plant process data are handled differently. Each type of data has different requirements for responsiveness, accessibility, and integration into services and tools within the NDP. This section discusses the requirements associated with data storage, thus providing insight into the considerations behind selecting storage types.

Enterprise record datasets provided by nuclear utilities are the primary foundational datasets within the NDP. This data includes data for components, stock, WO's, and maintenance plans as an example. These data sets are highly structured and require minimal manipulation or calculation before being displayed within the application, which suits storage within a relational database management system such as Azure SQL Database. Data associated with plant records are structured with constraints to ensure that the data entered are valid for a particular record. This is well suited to SQL databases that enforce data integrity.

Record data stored in SQL databases displayed within the application can be extracted, transformed, and loaded into a denormalized database using tools such as Azure's Data Factory service. Utilizing a denormalized database as an intermediary between the raw data and the application allows the raw record data to be stored separately from application-specific datasets. The denormalized database is also pre-joined and filtered to support search and analytic speed improvements. Database indexing and optimization are also utilized to improve the performance of data requests associated with the application.

Due to the data integrity requirements, big data analytics are not well-suited to being performed in SQL within the NDP's structured and constrained environment, thus, an alternate option for storing data is required. Azure Data Lake Storage Gen2 (ADLS) is an example of a storage technology well suited to analytics since it is a more flexible data storage service. PKMJ opted to use ADLS due to the flexible features associated with storage of various file types that support analytics on large, diverse datasets.

ADLS is a form of blob storage with a hierarchical folder structure in the form of containers, thus affording users additional convenience when organizing diverse file types. The best practice for configuring the ADLS is by utilizing three tiers for data: bronze, silver, and gold [7]. The bronze tier consists of raw data used for storing raw vibration and plant time-series data for use in the NDP. However, since the raw data is not ready to be processed, it must be manipulated in a separate service, such as Azure Databricks, to be utilized by CBM models. These transformations can be simple or complex depending on the intended use. The transformed data are stored in the silver tier and are ready to be consumed by models or analytic functions. The outputs of the models or analytics are stored in the gold tier and are ready to be presented to the end-user in front-end displays within the application or reporting tools. A diagram of the different data tiers is displayed below in Figure 5.

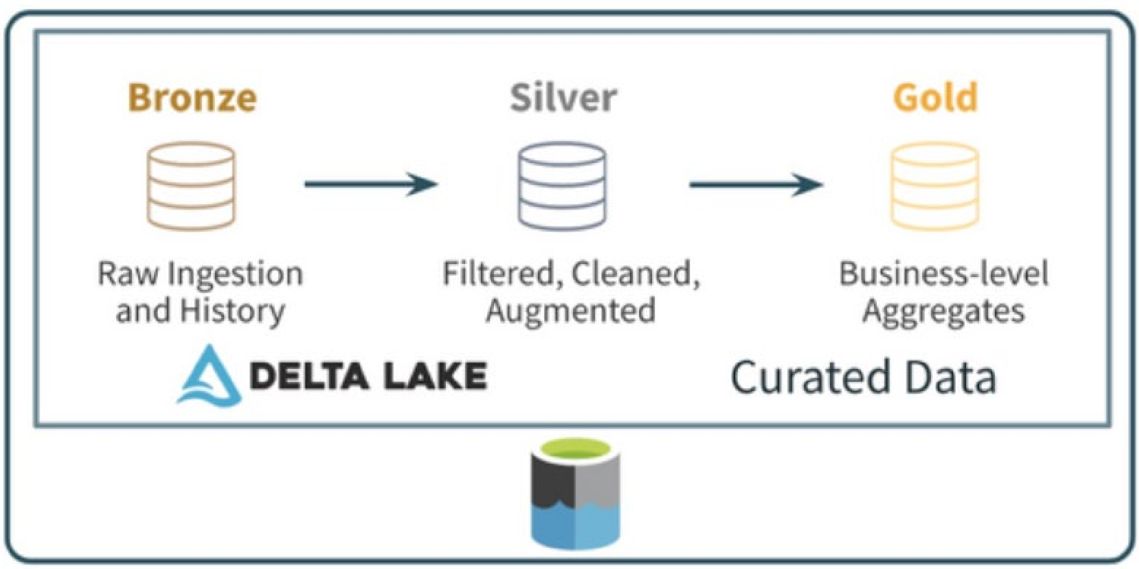

Figure 5 - Data Lake Data Tiers

Another specific function that the ADLS service improved under this project is the storage of raw vibration data. The vibration sensor vendor (KCF Technologies) utilizes Amazon Web Services (AWS) as their cloud provider of choice. As a result, the data storage configuration differed from the PKMJ NDP format (i.e., from Azure). The raw vibration data were not structured in a consumable format for PKMJ, as each entry in the vibration data table had 4,096 separate data values within a single field. For the vibration data to be consumable by the Fault Signature Model, the data had to be transformed to include a timestamp for each value adjusted for the sampling frequency of the associated vibration sensor. After the vibration datasets were transformed in Azure Databricks, they were stored in ADLS to support large scale analyses due to the size of the data.

Regarding cost, SQL databases are far more expensive to run compared to the ADLS service and should only be used when a full Relational Database Management System is required. Where possible, alternatives to SQL databases should be considered, with data integrity and constraints being ensured via methods other than the inherent functions associated with SQL databases. Data integrity checks can be performed via Databricks or other methods to ensure accuracy without creating SQL constraints and foreign keys. PKMJ's NDP originally considered the use of blob storage accounts as opposed to ADLS, as blob storage is a cheaper option per unit of storage. However, the hierarchical design of ADLS in addition to the functionalities unlocked by the ADLS service (as described above), provided benefits that justified the cost of the more expensive option. When aged-out data sets are no-longer needed by the application, stored data should be moved to archive-type storage on a scheduled basis to minimize cost and size of the SQL database.

Design of any digital platform in the industry should strongly consider the data requirements imposed on the platform. The security, flexibility, cost, agility, users, and use cases must be considered to determine the best option(s) for each platform. The NDP focused on utilizing technologies such as SQL databases and ADLS services to maximize the efficiencies of utilizing existing data structures. PKMJ's recommended practice is to evaluate data structures to obtain the optimal balance between cost and performance in terms of data storage and manipulation.

2.4 Build Deployment and Web Hosting

Another important consideration for the NDP was determining the appropriate application design framework. Many implications are associated with the application design, as the code must be configured to be functional for users while still being maintainable for developers. The primary options reviewed for the NDP included using a single page application design and a multi-page application design.

PKMJ opted to develop the NDP as a single page web application written primarily in JavaScript. The primary benefit of this approach was that it supported scaling the web application for additional utilities and consumers in the future. Single-page application architecture is suited to scaling out when additional resources are needed as opposed to multi-page application architectures which largely depend on scaling up. Scaling up refers to the use of a more powerful server to handle resource demands and scaling out refers to using multiple servers acting in parallel to scale beyond the capability of a single server. The Azure infrastructure supports this scaling out functionality with their resources, thus further supporting the selected approach. In terms of the NDP, this approach means that the addition of new sites, units, or industry customers is less likely to impact functionality or performance for existing customers and reduced hosting costs will be incurred as a result of the automatic scaling capabilities within Azure.

Single-page applications offer additional advantages in terms of presenting data to users over a single web space as most resources such as HTML, CSS files, and scripts are only loaded once from Azure Blob Storage. Furthermore, development is simplified thanks to not having to render individual pages within the application. The authorization within a single-page application is more complicated compared to multi-page applications, but the authorization framework using stateless API's is well-established and was suitable for the application being developed.

The collection of files that render the application are updated by the developers using SourceTree, which is a common Git Graphical User Interface (GUI) for repositories. To perform local development, the developers can load and clone the files that make up the web application out of the blob storage account. This capability is further simplified for a single-page application as the files that render the working application can be easily displayed. As the developers make changes to the application, the changes are overwritten within the Azure storage account.

The web hosting design for a single-page web application is consistent with the reference architecture developed by Microsoft for a serverless web application as it is well-suited for use in the Azure Cloud [8]. In this configuration, the term "serverless" implies that the application utilizes Backend as a Service (BaaS) as well as Functions as a

Service (FaaS) to ensure that developers do not need to directly deploy, configure, or manage servers (even though servers are used, the developers do not interact with them directly). Selecting this approach also allows for dynamic allocation of computing resources, consumption-based pricing, and scalability based on demand without performing manual configuration. In the scope of the NDP, this design allowed developers to focus on the page content without worrying about resource limitations due to automated scaling. Additional design requirements are required within the application to ensure that consumption is not excessive (e.g., preventing tasks running in an infinite loop), but this is addressed by implementing good industry development practices within the application.

PKMJ recommends using Continuous Integration/Continuous Deployment (CI/CD) within a digital platform software design to further enhance the availability of the application. CI/CD allows code changes to be deployed to the live page as soon as they are incorporated. This differs from traditional deployment, which requires downtime of the application to delete or modify application files. As the live application is always up to date, it minimizes some error traps that accompany changes being made within the application and provides additional insights into the application's current state of development.

This approach is well suited with integration into Microsoft's DevOps service as DevOps features deployment capabilities integrated into its work planning and repository tools. DevOps can support functionality such as automating tests for new and updated code to ensure that the code coverage meets anticipated standards based on the stage of development. DevOps can also manage the files being hosted, which further supports testing and development within an application. The synergies between DevOps as a service, the selected web hosting approach, and CI/CD are recommended to maximize efficiencies for platform development and deployment.

3. DIGITAL PLATFORM USER INTERFACE

The NDP is designed as a centralized solution tool for fostering effective decision-making and enhancing knowledge of plant condition. As discussed above, the selected format for this tool was a browser-based application that can provide information to the appropriate stakeholders in an easily digestible format, thus enabling decisions to be supported by the latest data. The data analysis techniques that support the NDP can be considered more complex than desired for dissemination to utility stakeholders, so the NDP must present complicated data analysis using simple visualizations and tables to convey the appropriate message to users.

3.1 Overview

PKMJ developed the NDP with a focus on providing nuclear power industry customers with a valuable tool to support and improve efficiencies within existing utility processes. The browser interface for the application provides focused visualizations that enables users to perform in-depth examinations of information associated with potential component alerts and opportunities from the supporting business intelligence tools. The NDP's initial design is meant to allow industry consumers to perform the following tasks digitally:

- Reduction of Unnecessary Plant Maintenance
- Alerting Site Personnel to Degrading Components via CBM Models
- Data Review and Analysis for Plant Records
- Automated Work Management, including considerations for Materials, Craft, and Scheduling (Reference Section 6)

Today, addressing the above tasks requires significant manual effort with extensive interaction among stakeholder groups (e.g., Operations, Maintenance, and Engineering), as well as accessing data sources in multiple locations. For example, addressing a plant alarm requires Operations to identify the alarm and implement immediate mitigating actions, Engineering to evaluate the required long-term corrective action, and Maintenance to plan, order parts, and perform emergent corrective maintenance. Identification of component degradation prior to notification via the plant alarm provides additional flexibility to determine a solution, minimize costs, and maintain the functionality of plant components. Development of a centralized user interface that integrates multiple data sources will improve customer processes by minimizing the interfaces required between stakeholder organizations and automating simple decisions associated with each task.

3.2 Maintenance Optimization

A major area for potential O&M cost savings is the reduction or elimination of unnecessary maintenance for plant components. The maintenance strategy utilized on a given plant component is based on vendor recommendations, industry feedback, or engineering review of the component's history. Enhancing this process can be accomplished by improving the engineer's access to both component data and historical WO data. PKMJ has developed a large database of component and maintenance historical data from nuclear stations throughout the world. With the knowledge and experience gained from examining the industry data, PKMJ collaborates with utility partners such as PSEG to review upcoming, scheduled maintenance activities and locate areas of high maintenance burden and large part costs associated with PM activities. In Phase 1 of this project, the maintenance optimization review was performed with significant intervention from data analysts to identify the appropriate maintenance targets and to identify the optimal value for a given customer [2]. Once the target scope was identified, the engineering review was performed using spreadsheets to organize the information. This section describes how the NDP application enhances this function and makes it more efficient.

The NDP allows users to review planned maintenance tasks by reviewing the scheduling dashboard within the application. This dashboard shows visualizations for WO's initiated, planned, or approved each week as well as insights associated with those WO's. In addition, visualizations of WO-related costs are also shown. The WO visualizations can be used to target planned work orders associated with Maintenance Plans to determine which WO's are candidates for work deferment up to two years in the future.

Maintenance plan data can be reviewed within the NDP by reviewing the specific plant records associated with each maintenance plan. The NDP allows users to sort the maintenance plan data based on parameters that include, but are not limited to maintenance due date, cost of parts associated with the maintenance, man hours for performing the work, and maintenance frequency. Using this functionality, a user can target PM tasks performed during refueling cycles, at various frequencies, based on cost, or a combination of other factors to quickly narrow down the scope of review. Flexibility in performing this search enables users to perform focused, efficient reviews of existing maintenance plan data to identify critical targets for evaluation.

Early in this project, the data sets were retrieved using PKMJ-developed algorithms to identify the PM tasks associated with the CWS that were relevant to the scope of work. Following approval of the scope, engineers reviewed the selected target maintenance tasks to determine whether any of the tasks could be optimized with a goal of reducing excessive maintenance on CWS components. This task was performed using access to the PSEG Enterprise Management System (EMS) to review the appropriate component and WO data. This review focused on determining the requirements and history of the components being reviewed. Once the appropriate data sets were identified, engineers completed their evaluations using spreadsheets that were then compiled and submitted to PSEG for approval. The recommendations made during Phase 1 of the project were accepted by PSEG for implementation. Table 1 lists the recommendations made during Phase 1.

Component	Task Title	Current Frequency	Industry Average	Recommendation	Recommended Frequency
Pump	**Refurbishment**	**6 years**	**14 years**	**Less Frequent**	**9 years**
	External visual inspection	18/24 months	2.8 years	Keep	18 months
Motor	**Vibration analysis**	**3 months**	**5.5 months**	**Less Frequent**	**6 months**
	Oil analysis	6 months	8 months	Keep	6 months
	Inspect/Electrical Testing	3 years	3 years	Keep	3 years
	Replace Motor	**6 years**	**10.7 years**	**Less Frequent**	**9 years**
Motor Cable	VLF TAN-Delta Testing	6 years	7 years	Keep	6 years
Protective Relays	Inspect/Calibrate	6 years	4 years	Keep	6 years
Pressure Switch	Calibration	4 years	4.2 years	Keep	4 years

Table 1 - PMO Recommended Frequency Results

The process of evaluating the PM strategy for each component and making the associated PM recommendations would be streamlined by using the newly developed NDP. Based on the evaluated targets, maintenance insights are generated within the NDP to allow users to evaluate each targeted maintenance plan for review. Maintenance insights are generated against existing maintenance plans as well as their associated components and WO's as applicable. Once generated, the user can drill down into the details of the associated maintenance insights to perform an engineering evaluation of the maintenance strategy for the components.

Data sets from the plant's EMS are displayed, partially automating the engineering justification by analyzing component and commitment data associated with the maintenance plan. Another integrated aspect is the review of component work history by displaying all historical WO's associated with the maintenance plan or component. Depending on the degree of integration with a utility's EMS, the amount of engineering work could be further reduced by addressing the various requirements examined when considering changes to a component's maintenance strategy. For fields that must be manually completed by an engineer, the user is provided several free text fields for entering engineering data to support or reject the recommended change.

Once the initial evaluation of the maintenance insight is performed, a reviewer can enter the NDP and confirm the results of the engineering work to determine whether the change is appropriate. The reviewer can leave comments for the preparer of the evaluation, and these comments can be resolved per normal Quality Assurance processes. Preparation and review of the maintenance insight is performed by PKMJ before being provided to the customer for their feedback and owner's acceptance within the NDP. Once all internal and customer comments are addressed, the maintenance insight evaluation can be marked "Complete" and implemented within the utility's EMS.

The above described maintenance optimization work saved Salem an estimated $4.37M over six years from 2021 to 2026 based on the limited scope associated with the CWS pumps and motors as described in the Technical Report for Goal 1 [2]. Future work scopes will be performed using the NDP and its associated toolset, enhancing the value of Maintenance Optimization services for utility customers such as PSEG. The integration of the NDP into the Maintenance Optimization process is an example case in which digitalization of a process can result in reduced maintenance costs for the utility while still maintaining component reliability.

3.3 Generation of Degradation Risk Insights

Equipment Reliability, or the reliability of plant components, is another key factor that highly correlates with NPP O&M costs. Components that are degrading in condition or reaching the end of its useful life begins to pose a higher risk of failure. Once components degrade, corrective maintenance work needs to be scheduled to restore the plant to an optimal state. These events also introduce additional Operational burden as failure scenarios require responses and monitoring from Operations personnel. The NDP is scoped to identify component and system health issues as soon as possible to provide utilities with the maximum amount of time to mitigate risks, plan corrective actions, and obtain supply chain support for work.

Today, utilities are often reactive in terms of identifying and responding to degradation risks. The degradation of plant components is typically identified via alarm indication aligned with process setpoints or limits. For example, the Salem CWS motors have a trip setpoint for when the motor outboard bearing temperature reaches a value that would damage the motor (185°F). Prior to reaching this limit, the plant's warning alarms (at 175°F and 180°F) provide Operators with an indication of the developing abnormal condition. Once the warning alarm actuates, Operators can mitigate the abnormal condition prior to reaching the trip limit. This design requires the implementation of multi-tiered alarms and alarm response procedures for each, necessitating significant effort from an alarm design and Operations standpoint. A similar approach is applied to most critical site parameters, but not every parameter features multi-tiered discrete alarms to the degree described above.

Artificial Intelligence and Machine Learning (AI/ML) models are tools that can be developed, along with a digital platform, to evaluate abnormal component conditions. With inclusion of the applicable instrumentation data, several AI/ML techniques can be applied to the data representing the components to determine whether conditions are normal or abnormal. A prime example of this is using AI/ML models to evaluate the condition of CWS components when input conditions are abnormal. The CWS may appear operate normally during an event, but this may be misleading masking from another abnormal system condition is involved. For example, if a pump is taken offline, a person might expect a different system response from the components remaining in service. If the system appears to be running normally with a pump train offline, the actual condition of the components may represent a secondary degradation. Such data analysis might be captured by experienced operators who understand the system response under abnormal conditions, but modeling techniques may identify the issue more reliably. Interpretation of plant data can be enhanced by using modeling techniques to discover conditions that may not be identifiable to humans as a result of the amount of data being evaluated or the minute nature of the deviations from normal states.

The NDP, as developed under this project, incorporates the modeling scheme developed by INL to show how integration of modeling into a digital platform can benefit the industry. The degradation risk insights are driven by plant process data, vibration sensor data, and plant design requirements. Combining these data sources for use in modeling instills confidence in understanding and responding to varying plant conditions. Examples of data used from a CBM tool or analytic model for incorporation into the NDP is presented in Table 2 below.

Fault Signature Output	
Field	*Comments*
Component ID	Component Identification – Allows alignment to Component/System Data from a Utility EMS
Date of Alarm	
Time of Alarm	
Condition(s) Identified	Failure Mode identified by the Model (ie. Diffuser Failure, Waterbox Fouling)
Sensor(s) tied to Alarm Condition	List of applicable sensors
State of Sensor(s)	Sensor data and simple statistics
Prognosis	Amount of Time in days before Condition Worsens or Component Fails
Maintenance Plan to Correct Condition	Mapping Failure Mode to Corrective Maintenance (Maintenance Plan or Historical Work Orders)
Immediate Threat to Public Health/Safety (Y/N)	Map to available diagnoses – Supports prioritization
Accelerated Degradation? (Y/N)	Map to available diagnoses – Supports prioritization
Component Object Type	(ie. Pump, Motor)
Subcomponent Impacted	(ie. Bearing, Lubrication, Casing, etc.)
Symptom Identified	(ie. Vibration, High Temperature)
Degradation Mechanism	If available (ie. Loose Hardware, Blockage, Wear, Leakage, Corrosion, Fouling)
Degradation Influence	If available (ie. Debris, Moisture, Contamination, Structural Resonance, Fatigue)
Risk Levels/Probabilities	Internal Processing for Risk Tolerance
Automated Work Management	
Field	*Comments*
Associated Work Solution	Maintenance Plan or Work Order associated with the diagnosis
Solution Data and Details	(ie. Hours, Tasks, Materials, Labor Type, etc. as available)

Table 2 - Example Condition Based Monitoring Outputs for Digital Platforms

Several modeling applications running in parallel can be utilized by the NDP if model outputs are manipulated to a format that can be processed by the platform. The NDP supports most standard data formats, thus, the platform is designed to be agnostic to the source(s) supplying the model. Utilizing multiple models or techniques is a best practice when understanding degradation risk identification capabilities within a digital platform as modeling using a single technique or technology stack may not accurately evaluate all possible events or component conditions. Modeling techniques of varying strengths and weaknesses can be utilized together to produce a holistic view of Equipment or Component Reliability. Utilities who have developed their own models as part of monitoring centers could combine their results with those of third-party vendors who possess mature model capabilities to provide a robust degradation risk alarm system.

3.4 Viewing Enterprise Management System Data

Utilities maintain their data records in their Enterprise Management Systems in support of plant functions. An EMS provides an auditable record of plant configuration which is critical for ensuring compliance to site design and licensing requirements. In addition, an EMS serves as a centralized location to support decision making by personnel reviewing data associated with their respective plants. The data required for the above-described use cases are incorporated into the NDP, eliminating the need to access multiple data sources. As personnel look to provide services associated with the NDP, having access to the EMS data is critical for efficient, accurate decision making.

Data to be made available in support of the desired functions should be adjusted based on the scope of services provided by the associated digital platform. For the NDP developed in this project, the target services provided are focused on Preventive Maintenance Optimization, Automated Work Management, and Degradation Risk Identification. Based on these target goals, the primary data made available within the NDP included component data, WO data, maintenance plan data, and stock data.

Component data sets are foundational for understanding the current plant configuration and its associated requirements. The scope of component data includes component-specific data fields such as component ID, associated system, manufacturer, model, and quality class. These data fields, along with others provided within the NDP, provides the basis for decisions associated with how components fit into plant design requirements for maintenance strategy evaluations and enable clear identification of components associated with alarms generated by the system. Without these data fields, the digital platform would lack detailed information on the component or equipment subject to the relevant services provided in the NDP.

Work Order (WO) data comprises another critical dataset for the NDP, as historical and planned work performed on components are used as the basis for decisions regarding plant components. For example, knowing the WO history for components can aid in decisions regarding maintenance effectiveness. In addition, WO data can be used to plan future work by utilizing historical work as a benchmark or template for creating automated WO's generated by the NDP. These functions mirror real-life actions taken by engineers and planners to streamline and improve their work efficiencies and make appropriate decisions.

Maintenance plan data comprises the target dataset being modified when performing PMO services. Understanding the current maintenance strategy for various components provides utilities with insights into work planning, as they can coordinate their maintenance work in specific time windows to minimize component downtime. Another insight gained from maintenance plan data is an understanding of the typical strategy for maintaining components of a selected type. These insights can be used along with data from vendors, manufacturers, historical sources, and the industry to inform maintenance strategy changes. Maintenance plan data are also a potential input to automated WO's, as the corrective action for resolving a degradation risk may be to implement the associated maintenance plan.

Stock data sets are used as an input for automated WO's as the piece parts required to perform maintenance plans or WO's are included as part of the NDP. The stock data includes items such as the total stock within the warehouse, total items on order (or due-in), and average unit price. When the stock required to perform a given task is understood, the site can perform the supply chain services required to obtain the associated parts. The stock data also forms a basis for evaluating the parts cost for performing work, which is a critical point in evaluating which maintenance tasks to focus on when making maintenance strategy changes. High-cost, low-value tasks are ideal targets for PMO activities.

4. CONDITION BASED MAINTENANCE MODELS

The ability to make notify plant personnel of possible degradation of components based on CBM modeling remains a key goal of this project. CBM models generate valuable insights for utility stakeholders by providing near real-time updates on plant conditions, while using advanced data analytics and predictive modeling. CBM models require supporting analytics to identify problems with plant components, specify with confidence the nature of those problems, and provide insight into the time until component degradation.

As the industry continues to examine cost efficiencies for O&M of their respective plants, it is critical to include CBM as an element of an overall strategy for maintaining plant health. This section describes the research performed to demonstrate CBM modeling for the Salem CWS components. The models use historical data to identify representative fault conditions for the components, employing detailed data analysis of the available process and monitoring data sources. The CBM modeling framework developed for the Salem CWS is a reference usable to expand the framework to various systems and utilities within the industry.

4.1 Architecture

The approach utilized for CBM in this project is a multi-stage modeling process with several interfaces to data sources of various types. Utilization of a multi-stage model was decided upon for research purposes, enabling specific examination of every element of the model to determine the benefits and drawbacks of each. This approach ensures modularity between elements of the model, but also introduced new complexity in getting each model element to interface with each other. A diagram showing the elements of the model is shown in Figure 6.

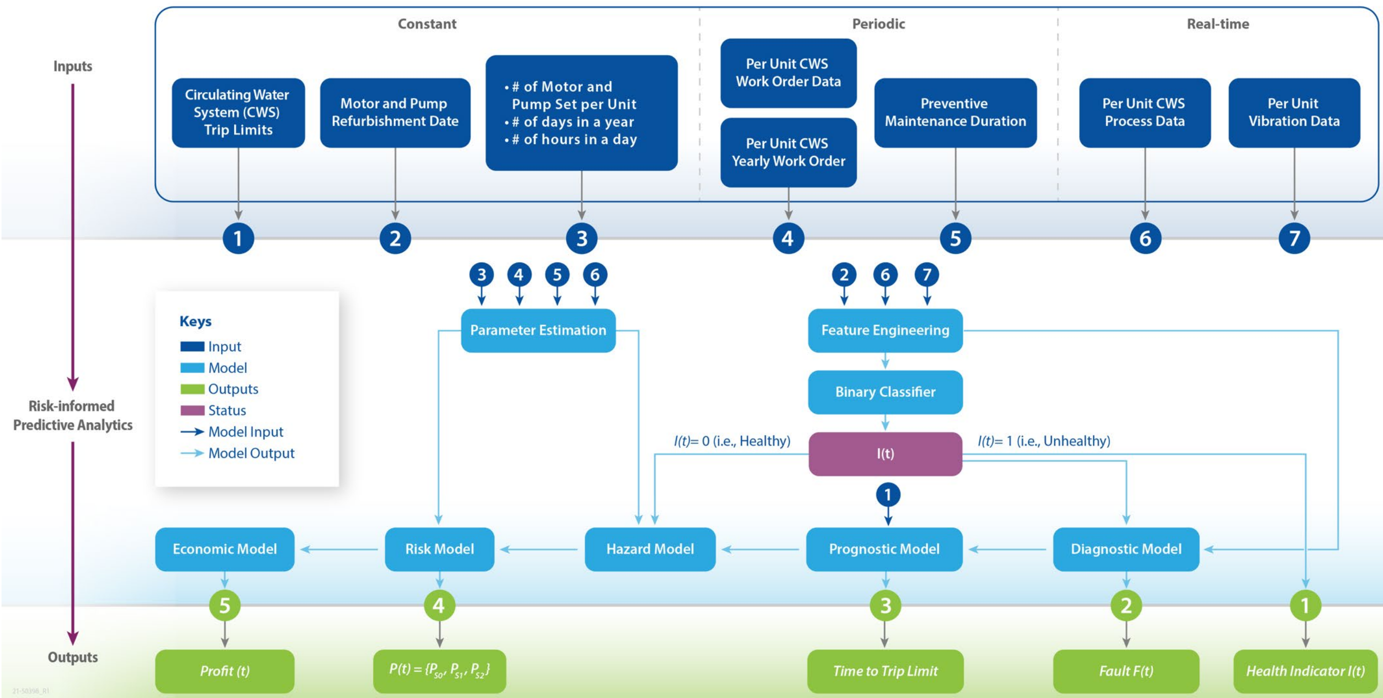

Figure 6 - Risk-Informed Condition-Based Maintenance Framework

To run the various models in the NDP, each modeling element must be configured to align with the corresponding inputs and outputs. PKMJ opted to use Azure Databricks as the environment in which to perform the modeling work as Databricks is well suited for performing large data analytics. The primary data inputs to the model are WO data, plant process data, and vibration data. These data inputs are inputted into Databricks to allow for processing within the model.

The best option for processing the data was to place the corresponding input files in Azure Data Lake Storage (ADLS) as bronze (raw) data as described in Section 2.3 and perform the processing in Azure Databricks. The structured inputs (e.g., enterprise records associated with maintenance plans and WO's) are stored in SQL databases. These data sets are then extracted from the corresponding database to a blob storage table within ADLS to perform the analytics. The plant process and vibration sensor data are time-series inputs stored in a blob storage account within ADLS. As part of this project, some data inputs such as plant setpoints, motor and pump refurbishment dates, and plant configuration data were stored in separate spreadsheets which were converted into tables for entry into ADLS and subsequent use with the model.

The data inputs in their raw formats are not ready for ingestion into the model, so pre-processing is performed to format the inputs for utilization within the model. Generally, this entails making the data consistent with the model-requested code by assigning timestamps to process data points to ensure they can be ingested by the model, reconfiguring tables to align with the intended data structure, and referencing storage locations for data instead of spreadsheets. Once the data sets are up to date and the references are aligned with the environment, the data sets are moved into the ADLS silver tier and the model can begin processing of the data.

As the models create intermediate output files, those files are stored in the ADLS silver tier so the next model can retrieve and utilize the files as input. These files do not represent the final output of the model as additional analysis is required, but they do represent enhanced features for the existing data. This configuration also allows for manual or automated checking of intermediate outputs to ensure that each of the model's outputs are configured as expected and that the outputs are reasonable based on the inputs. An important best practice for a complex modeling process is to ensure that error checking and logging features are implemented whenever feasible to support troubleshooting.

The final model outputs are placed in ADLS gold tier to represent the consumer-ready outputs of the model. These outputs are formatted to support implementation in the NDP front-end application or provided to a reporting and visualization tool. The final outputs typically take the form of a spreadsheet; however, this form can vary, depending on the model variables. Thus, use of ADLS or blob storage is well suited to enable more flexible analytic outputs for later consumption. Additional details on how the front-end utilizes the model output are included in Section 6.

4.2 Fault Signature Analysis

A fault signature is a unique set of fault features used to assess the current and future health condition of the plant asset which is undergoing a particular fault. One of the aims of the fault signature is to enable informed decision-making to prevent potential failure of the plant asset. The fault signature can also be used for root cause analysis if failure occurs. The different fault modes associated with a plant asset have unique, identifiable fault signatures. In practice, fault signature identification and diagnosis is not straightforward and can benefit from historical data. Each detected fault signature for a degradation mode should have enhanced feature verification and confidence by selecting additional process and condition monitoring data that provides complimentary information.

The target CWS faults that were reviewed as part of this project include:

- Waterbox Fouling
- CWP Diffuser Failure
- CWP Bellmouth Failure
- CWP Shaft Misalignment
- Clogging in Air Intake Screens of the CW Motors
- Moisture and Salt Contamination of CW Motor Windings
- CWP Low Oil Level

Based on the period encompassing the data made available for analysis (i.e., 2008-2021), few fault types have resulted in multiple plant derates and trips (i.e., impacting plant generation), with others impacting plant generation once or not at all. In this report, the plant process data, vibration sensor data, and WO data associated with the CWS are used to develop a CBM modeling solution. The CWS WO data [1] are used to create an approximate timeline of when faults occurred and were corrected, in addition to a timeline of PM activities. In particular, the fault timeline is important for identifying possible features relevant to enlisted fault modes that, in turn, are used as training data sets for ML algorithms.

The process of obtaining the training set data via this research effort consists of the following steps. The WO database or a utility's EMS notifications are searched for CWS faults of interest. The plant process data associated with CWS during the time when the fault occurred is manually searched to uncover the corresponding time when a Circulating Water Pump (CWP) was taken offline. The process and condition monitoring data chosen for the ML training sets are extracted from the time history by tracking backwards in time to when the CWP was taken offline for repair. A subset of this extracted data is then used to select the fault signatures. The importance of determining when the CWP was taken offline is based on two major assumptions: (1) the data history prior to the identified CWP shutdown contains the desired fault signatures, and (2) the data history following the CWP shutdown represents normal CWP operation.

Of the faults of interest examined during this project, only four fault types had multiple instances that caused CWP shutdowns. The multiple occurrences of these faults strengthen the fault signature analysis and its usage to train/test ML algorithms. Waterbox fouling has numerous instances of causing CWP shutdowns. While the waterbox fouling fault is not a pump or motor fault, it is a system fault that may show symptoms affecting performance of the pump. The waterbox fault also occurred numerous times, allowing for development and testing of CBM algorithms. Fault types with only a single instance of causing the shutdown of a CWP provided limited information for developing a fault signature and training ML algorithms. These potential fault signatures contained within the data are not readily resolvable at this time. It is anticipated that, as the ML technology matures for operational plant applications, these subtle faults will be identified.

The example of fault signature for "Clogging in Air Intake Screens of the CW Motors" affords a good understanding of the procedure for selecting the plant process data and identifying the progression of the fault. In the years since the single fault occurred, the plant has done an excellent job of preventing it from progressing into significant degradation. It is anticipated that CBM will provide obvious features/outputs that will be as successful in preventing asset degradation from other faults. The air intake fault and approximate date of the shutdown occurred were found by searching the WO database.

Figure 7 shows the small-time window from the historical process data for a CWP motor affected by the clogged air intake. This time window contains the fault timeline reported by the WO. In late May 2008, the CWP was taken offline before the clogging of the air intake. Prior to this the stator temperature was nominally 165°F during most of May. Following the May 2008 shutdown of the CWP, the stator temperature started out nominally at 165°F, then increased significantly to 190°F and continued to rise. In mid-July 2008, the CWP was taken offline to address the fault. Figure 7 indicates that the pump fault was corrected, since the stator temperature dropped back down to a nominal temperature of 160°F after the CWP is brought back online later in July 2008.

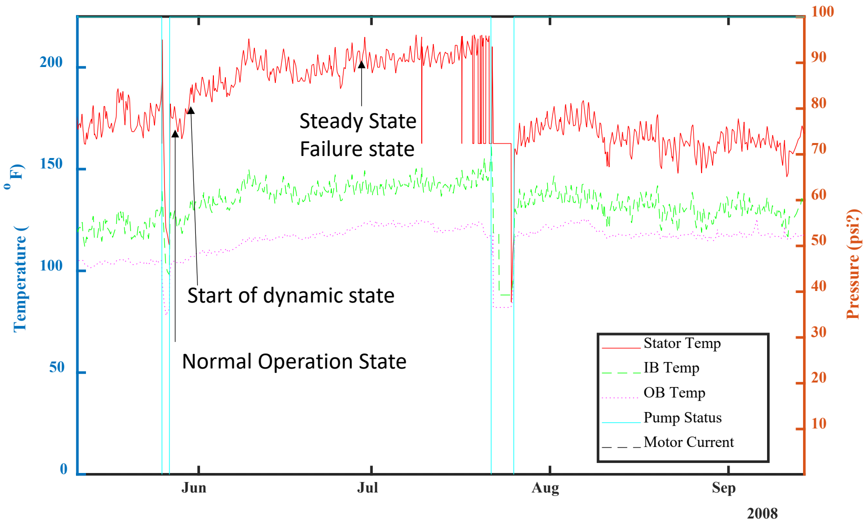

Figure 7 - The CW Motor Air Intake Clogging Fault Shows an Idealistic Progression of the Fault from Initiation to Correction

4.2.1 Feature Engineering

This section provides a brief discussion on features extraction from the CWS associated plant process data and vibration data as part of the fault signature development process. For more details on feature extraction, see [1].

4.2.1.1 Features from plant process data

For the CWS, the following plant process data along with other information are used as features, denoted as $X_{Features}$, are extracted for each M&P set:

- The *differential temperature (DT)* is calculated as the difference between the outlet water temperature (associated with the combination of each M&P set) and the inlet river temperature
- The measured *motor inboard (MIB) temperature*, *motor outboard (MOB) temperature*, and *motor stator temperature*
- The *motor current* measurement (available after September 2017)
- From the reset and targets files, historical CW M&P replacement/refurbishment dates are available; the M&P run-hours from one replacement to the next is considered in the calculation of the *motor age* and *pump age*
- To consider the seasonal effect on the data, *week of the year* is calculated for every timestamp and used as a feature.

Multiple features were extracted from the CWS plant process data to train and test the ML models.

4.2.1.2 Features from vibration data

The features extracted from vibration data for each M&P set are categorized into time-domain features and frequency-domain features. The extracted feature is a result of the vector summation of the X- and Y- directions for the locations MIB, MOB, motor axial (MA), valve axial (VA), and valve radial (VR) [1].

- *Time-domain features*: The resultant (i.e., vector summation of X- and Y- directions) standard deviation (σ_{VibXY}), and resultant mean ($\overline{X}_{VibXY}$) are extracted as time-domain features.

  Thus, a total of 10 time-domain features are extracted for each M&P set. In the plots and figures generated, each time-domain feature is represented in the format {location}_STD for resultant standard deviation, and $\{location\}\_Mean$ for resultant mean.

- *Frequency-domain features*: The fundamental frequency (1x) of the CWP motor is 4.97 Hz. The CWP has four vanes and six diffuser vanes. The frequency components are extracted using the fast Fourier transform. The band magnitudes of resultants (vector summation of X- and Y- directions) at the following harmonics are extracted:
  - Fundamental frequency (motor running frequency) 1x
  - Second harmonic of the fundamental frequency, 2x
  - Vane pass frequency ($\#\ of\ pump\ vanes * fundamental\ frequency$) (VPF)
  - Diffuser pass frequency ($\#\ of\ pump\ diffusers * fundamental\ frequency$)
  - Second harmonic of the VPF, 2 × VPF
  - Third harmonic of the VPF, 3 × VPF
  - Fourth harmonic of the VPF, 4 × VPF
  - Vane diffuser pass frequency ($\#\ of\ pump\ vanes * \#\ of\ pump\ diffusers * fundamental\ frequency$)
  - Electric line frequency at 120 Hz (2x LF)

A bandwidth of 1 Hz is maintained on either side of harmonics, except for 1x, 2x, and 2x LF. For 1x and 2x LF, the bandwidth of 0.1Hz is selected, and for 2x, the bandwidth of 0.2Hz is selected. The bandwidth at each harmonic frequency is selected by starting with the bandwidth of 0.1 Hz and increasing until no change in the total magnitude trend is observed. Hence, the bandwidths are different for each frequency of interest. There are a total of nine frequency-domain features at each location (such as MA, MIB, MOB, VA, and VR) of the M&P set, meaning a total of 45 frequency-domain features from the vibration sensor nodes for each Salem CWP. The frequency domain features are collectively represented as $M_{VibXY,freq}$ (see Figure 11).

4.3 Diagnostic Model

The diagnostic model was developed to estimate the condition of the CWS M&P set, based on the features extracted from the vibration and plant process data. A two-level cascaded diagnostic (one to detect the CWP's condition using binary classification, and another to detect faults using multiclass classification) was also developed. Both models were trained by combining all the CWS M&P-set-related features extracted for each fault and healthy scenario from Salem Units 1 and 2. A diagnostic model can be defined as $Y = F(x, \theta)$, where F is the eXtreme Gradient Boosting (XGBoost) [9] (binary/multiclass) model, Y is the output of the model, x is the input, and θ denotes model hyperparameters. For XGBoost, the main hyperparameters [9] are $\#estimators$, $maximum\ depth$, $minimum\ child\ weight$, $column\ samples\ per\ decision\ tree$, and $gamma, \gamma$. Hyperparameter optimization finds a list of hyperparameters and associated values that yield the optimal XGBoost model. For both the binary and multiclass XGBoost models, the hyperparameter tuning was done using the *hyperopt package*, which is a combination of grid search and random search approaches (this package is part of the python ML library).

4.3.1 Binary Classifier for CWP Condition Diagnostic

CWP condition diagnostic is a binary classifier to estimate whether a given CWP is healthy or unhealthy. All the fault categories are grouped into a single category called *unhealthy*. For the binary classifier, the input includes all the features extracted from plant process data, and all the time-domain features from vibration data. The trained binary XGBoost classifier outputs the assessment, $I(t)_{ij}$ (healthy or unhealthy), the probability of assessment ($P\big(I(t)_{ij}\big)$), and the timestamp of prediction ($t_{degrade}$), as shown in Figure 8, where the subscript $\mathrm{i} = \{1,2\}$ is the plant unit, $\mathrm{j} = \{1A,\ 1B, 2A, 2B, 3A, 3B\}$ is the M&P set, and $\mathrm{k} = \{MIB,\ MOB,\ MA,\ VA,\ VR\}$ is the location of vibration sensors on the M&P set. The outputs of the binary classifier include, $t_{degrade}$, the time instant at which the pump or motor is determined to be in an unhealthy state, and $\mathrm{P}\big(\mathrm{I(t)_{ij}}\big)$ provides the level of confidence in estimating the state, denoted as $\mathrm{I(t)_{ij}}$.

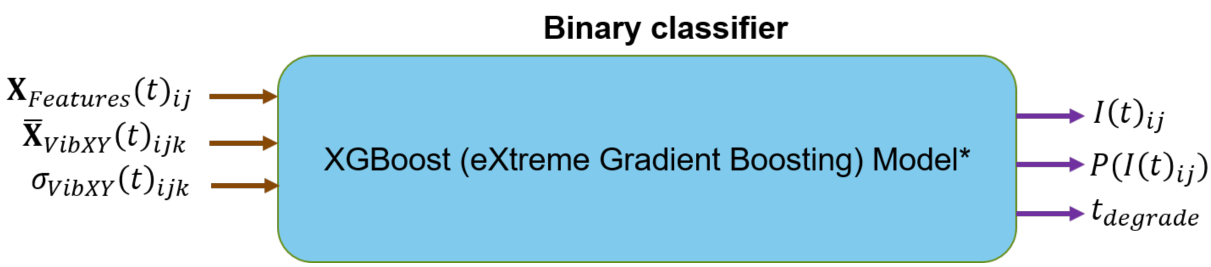

Figure 8 - Binary Classifier Model Using XGBoost to Predict CWS Components as *healthy* or *unhealthy*

The binary classifier model is evaluated for $F1$ score defined as $2 * \frac{precision*recall}{precision+recall}$ [10], where precision is defined as $\frac{TP}{TP+FP}$, and recall is defined by $\frac{TP}{TP+FP}$. Recall is also termed as “sensitivity” or “true positive rate”, and false positive rate is defined by $(1 - \text{specificity}) = \frac{TN}{TP+FN}$. Here TN, TP, FN, FP are true negative, true positive, false negative, and false positive, respectively. In assessing the CWP condition, $I(t) = -1$ corresponds to the *unhealthy state* and $I(t) = +1$ corresponds to the *healthy state.* The performance of the binary classifier is graphically visualized via the receiver operating characteristic (ROC) curve, using true positive and false positive rates. Then, the aggregated measure of performance across all possible classification thresholds is estimated using area under ROC curve (AUC) [10].

From Salem Unit 1 and Salem Unit 2 combined, there are a total of 299,457 and 53,300 *healthy* and *unhealthy* samples, respectively. Since, the number of *healthy* samples is significantly higher than that of the *unhealthy* class, the majority class, *healthy,* is downsampled to the size of the minority class, *unhealthy*. First, for the *healthy* class, only data after 2015 are considered. This also helps to be within the six-year CWP replacement dates, thus capturing the aging patterns of all the CWPs. Next, the *healthy* class is further downsampled using random sampling without replacement. For training the model, the data from 2008 to 2018 (the *healthy* samples are from 2015, the unhealthy samples are from 2008) are considered as training samples, the samples from 2019 till February 2020 as a validation set, and the samples from March 2020 till January 2021 as a test set. Note that the vibration data features are extracted only for the CWP diffuser issue, which occurred between April 1, 2020 and April 21, 2020. Hence, in the training and validation set, no vibration data were involved.

The performance of the binary XGBoost model for CWP condition prediction is shown in Table 3. In both the training and validation cases, the recall and AUC score surpass 0.8 in predicting *unhealthy* class. But for the test cases, the recall is around 0.63, clearly indicating a large FP prediction. The overall AUC curve for training, test, and validation is shown in Figure 9. The results also indicate that the model is too sensitive when predicting the *unhealthy* class. The main reason for the high FP rate is the incomplete data upon which the model is trained (for example, motor current is available only after September 2017). Also, for certain considered faults such as *misalignment*, a particular fault signature from the plant process data was not seen, thus contributing to false detections. Besides, the fact that baselines for data change after major maintenance can also influence accuracy. In Figure 10, it is evident that temperature distributions for fault *misalignment* aligns with the distribution of *healthy* data.

Data	Condition	Accuracy (%)	Precision	Recall
Training	Healthy	95.15	0.94	0.95
	Unhealthy	96.23	0.96	0.95
Validation	Healthy	91.00	0.85	0.97
	Unhealthy	89.00	0.96	0.83
Test	Healthy	92.17	0.85	1.00
	Unhealthy	77.42	1.00	0.63

Table 3 - Binary XGBoost Classifier Results for CWP Condition Prediction.

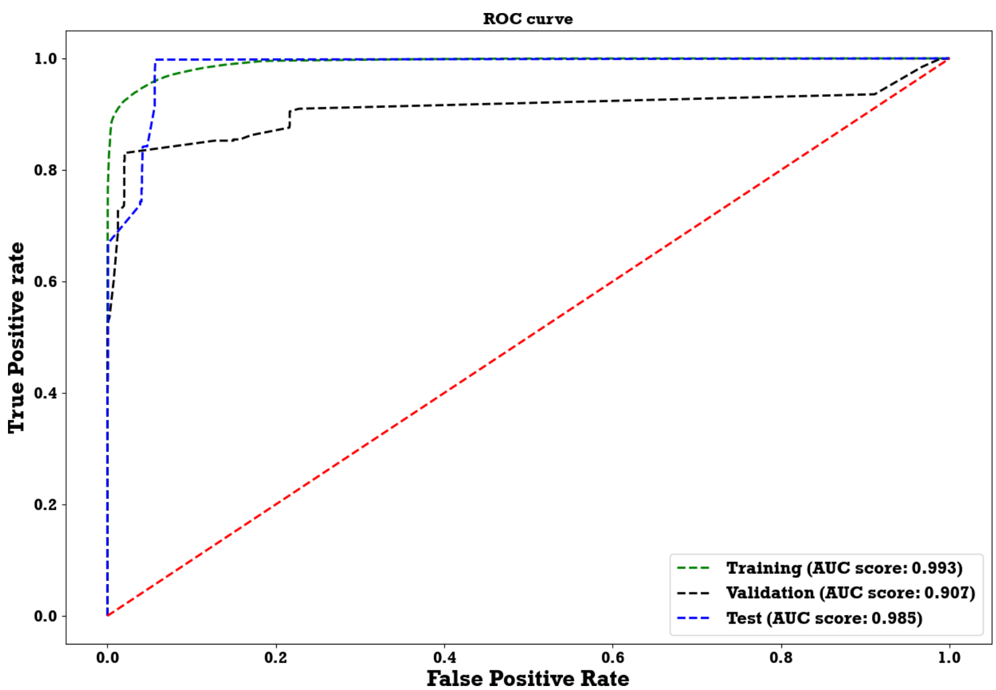

Figure 9 - ROC Curve for Binary XGBoost Classifier to Predict CWP Condition

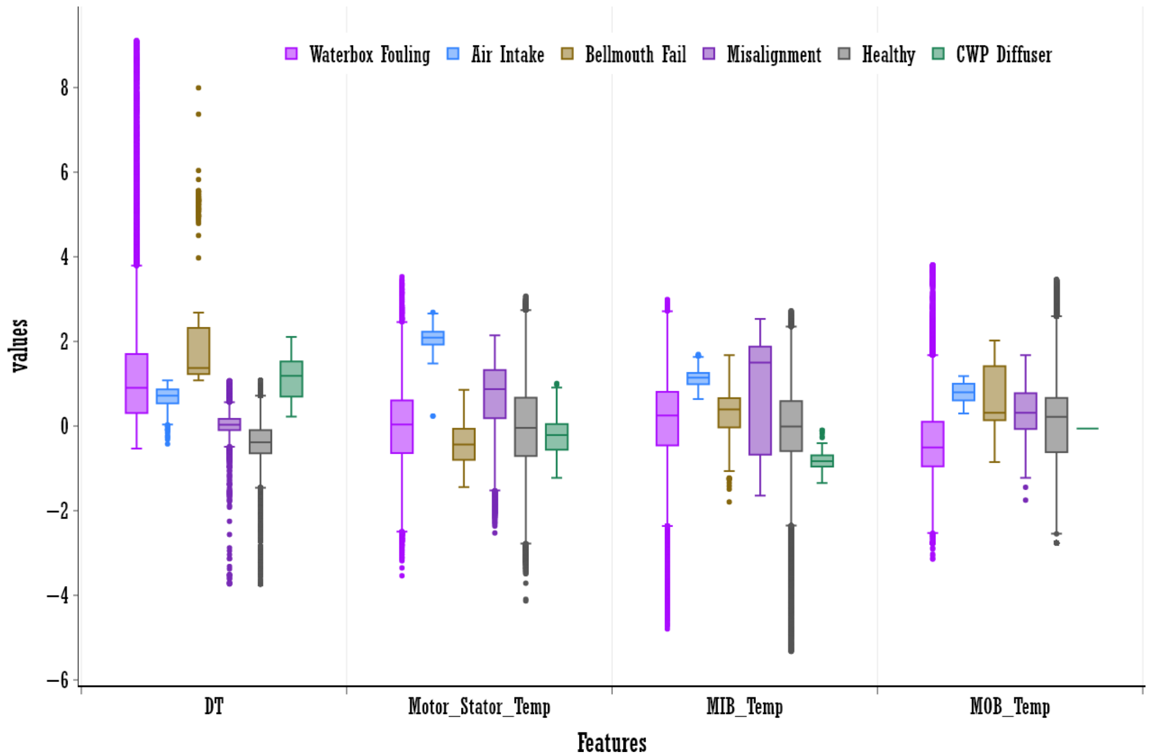

Figure 10 - Plant Process Data Distribution Across Different Faults. Only Motor Temperatures and Delta T are Considered as they are Consistent from 2008 to 2020

4.3.2 Multiclass Classifier for Fault Estimation

CWP fault prediction, a multiclass classifier based on the XGBoost algorithm, predicts a particular fault given that an *unhealthy* condition is predicted in the CWP condition diagnostic model. The target faults mentioned in Section 4.2 are considered, the associated data are extracted, and a multiclass classifier to estimate the underlying fault when the CWP is in an *unhealthy* condition is developed. For the multiclass classifier, the input includes all the features extracted from the plant process data (except *week of the year, pump age*, and *motor age*), as well as the frequency-domain features from the vibration data. At time instance *t* for CWP_{ij}, the trained multiclass XGBoost classifier outputs the estimated fault, denoted as, $F(t)_{ij}$, the probability of a particular fault estimation, denoted as $P\left(F(t)_{ij}\right)$, and the *n* is the number of most significant features $F_{features}(n,t)_{ij}$, as shown in Figure 11.

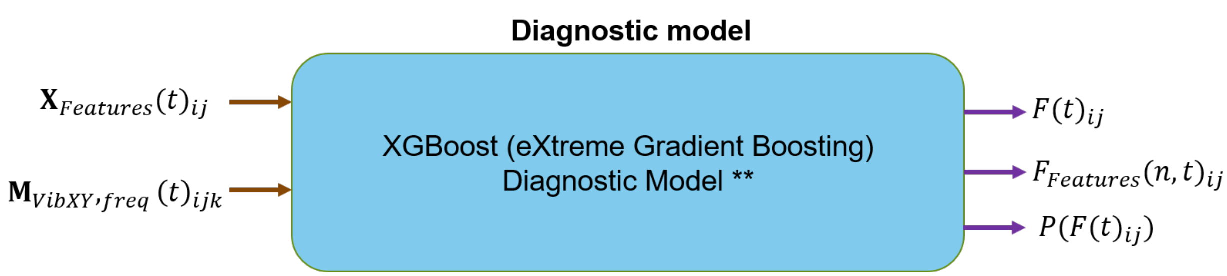

Figure 11 - CWP Fault Prediction Model Using Multiclass XGBoost Classifier

For the multiclass classifier, the fault data extracted for *waterbox fouling, CWP diffuser, misalignment, bellmouth fail,* and *air intake* totaled 46,289, 502, 5,237, 417, and 855, respectively. The *waterbox fouling* fault had the most data points. Thus, the *waterbox fouling* fault data are considered only after 2018, when the motor current data are also available, reducing the data points to 7,147. A full description of the data points split among the training, validation, and test sets for multiclass fault estimation is given in Table 4. Test samples were not used either during training or validation. Also, for cases like *CWP diffuser*, since there is only one incident, the samples from the last part of the degradation before maintenance are used for testing.

Faults	Training set	Validation set	Test set
Air Intake	**Year: 2008** **First 599 data points of the fault**	**Year: 2008** **208 data points (excluded from training and test set)**	**Year: 2008** **Last 48 data points of the fault**
Bellmouth Fail	**Year: 2010-2012** **#Data points: 249**	**Year: 2013** **#Data points: 157**	**Year: 2017** **#Data points: 11**
CWP Diffuser	**Year: 2020** **First 336 data points of the fault**	**Year: 2020** **108 data points (excluded from training and test set)**	**Year: 2020** **Last 46 data points of the fault**
Misalignment	**Year: 2013** **#Data points: 3618**	**Year: 2014** **First 1547 data points of the fault**	**Year: 2014** **Last 72 data points of the fault**
Waterbox Fouling	**Year: 2018, and January-September 2020** **#Data points: 4383**	**Year: 2019** **#Data points: 2173**	**Year: October 2020-January 2021** **#Data points: 591**

Table 4 - Data Split Among the Training, Validation, and Test Sets for Multiclass Fault Estimation

The multiclass XGBoost model's CWP fault prediction performance in terms of the confusion matrix for training, validation, and test sets are shown in Table 5, Table 6, and Table 7, respectively. Overall, a fault estimation accuracy above 99% was achieved for training, validation, and test sets in predicting the faults. In the tables, *actual* corresponds to the extracted fault data from plant process and vibration data (discussed in Section 4.2), and *predicted* corresponds to the fault detected by the fault estimation model. A few misclassifications are seen in which the *air intake* fault is estimated as *misalignment*, and *bellmouth failure* is estimated as *waterbox fouling*.

predicted→ *actual↓*	Air Intake	Bellmouth Fail	CWP Diffuser	Misalignment	Waterbox Fouling
Air Intake	588	0	0	11	0
Bellmouth Fail	0	247	0	2	0
CWP Diffuser	0	0	336	0	0
Misalignment	0	0	0	3618	0
Waterbox Fouling	0	0	0	0	4383

Table 5 - Confusion Matrix for Multiclass XGBoost Fault Prediction on Training Data

predicted→ *actual↓*	Air Intake	Bellmouth Fail	CWP Diffuser	Misalignment	Waterbox Fouling
Air Intake	198	0	0	10	0
Bellmouth Fail	0	152	0	0	5
CWP Diffuser	0	0	108	0	0
Misalignment	0	0	0	1547	0
Waterbox Fouling	0	0	0	0	2173

Table 6 - Confusion Matrix for Multiclass XGBoost Fault Prediction on Validation Data

predicted→ *actual↓*	Air Intake	Bellmouth Fail	CWP Diffuser	Misalignment	Waterbox Fouling
Air Intake	47	0	0	1	0
Bellmouth Fail	0	10	0	0	1
CWP Diffuser	0	0	46	0	0
Misalignment	0	0	0	72	0
Waterbox Fouling	0	0	0	0	591

Table 7 - Confusion Matrix for Multiclass XGBoost Fault Prediction on Test Data

4.3.3 Feature Significance

For the faults estimated using the fault signature model, the Shapley additive explanations (SHAP) [11] value is used to determine most significant features driving the estimation of each fault. Using the SHAP value, the most significant feature for each fault is shown from Figure 12 to Figure 16. Note, most of these estimations do not take into consideration vibration data (except *CWP diffuser*) at this stage of the research. The inclusion of vibration-based features can change the values presented in the figures.

In each figure, only the features used in the fault estimation model which have an influence value greater than zero are shown. Each bar represents the amount of influence each feature has on a particular fault estimation, and the accumulation of all the feature influences sum to 1. Figure 15 is of particular interest here. Based on historical plant process data assessment, no evidence of *misalignment* was found. As a result, developing and verifying feature significance based on SHAP values for the *misalignment* fault was challenging. For now, based on the algorithm approach followed, the SHAP values for the *misalignment* (yet to be verified) are presented in Figure 15. In the future, via simulation or based on the information obtained for a *misalignment* fault from other plant sites, these SHAP values would be updated.

From the identified significant features, the most dominating feature is forecasted to observe the trend of an identified fault. If the most significant feature is not available, then the next dominant feature can be used for forecasting. For example, in *CWP diffuser* fault, if the vibration data is not available then *MIB_Temp* or *Motor Current* can be used for forecasting.

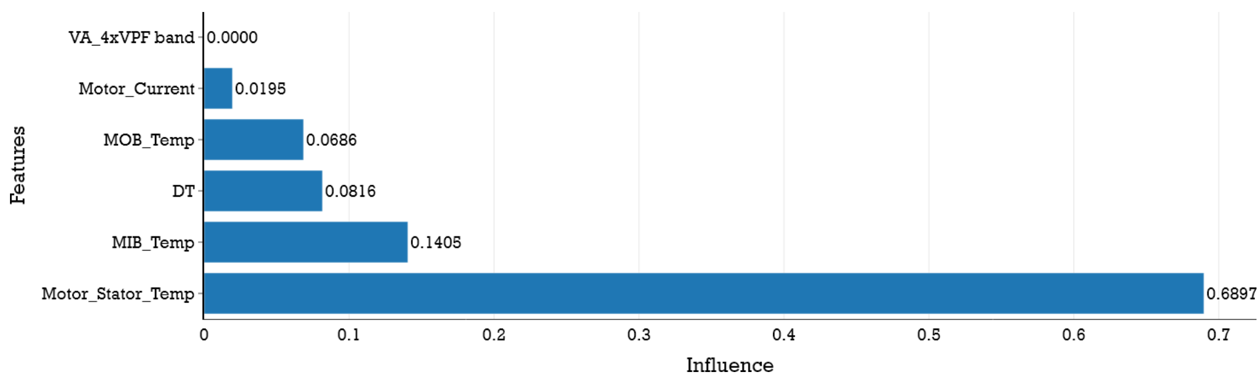

Figure 12 - Feature Influence on Predicting *Air Intake*

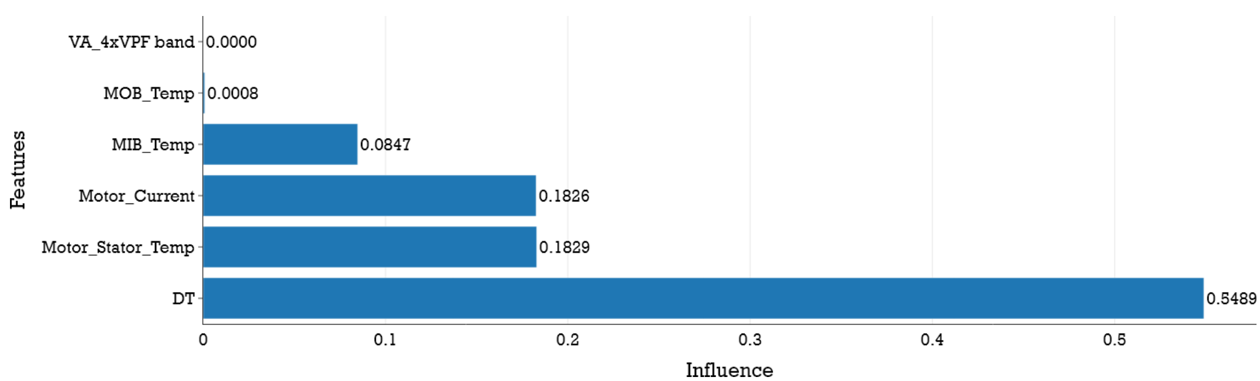

Figure 13 - Feature Influence on Predicting *Bellmouth Failure*

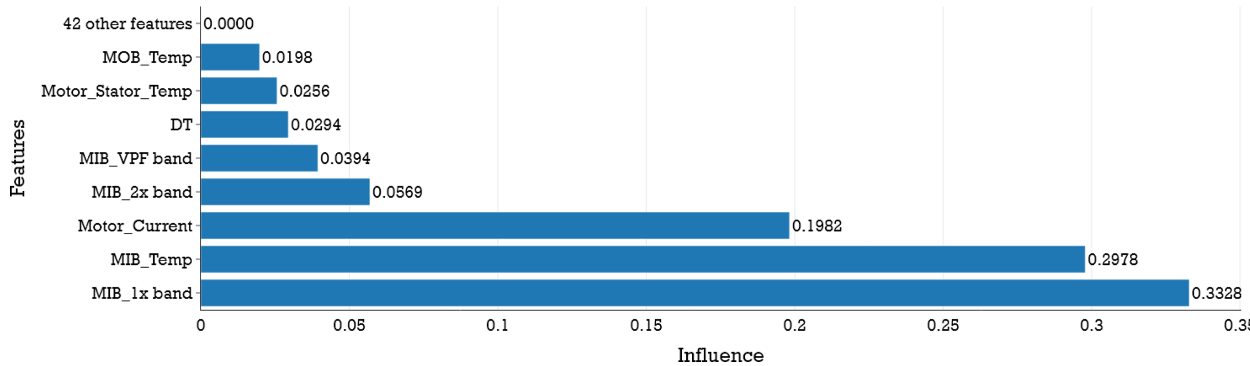

Figure 14 - Feature Influence on Predicting *CWP diffuser*

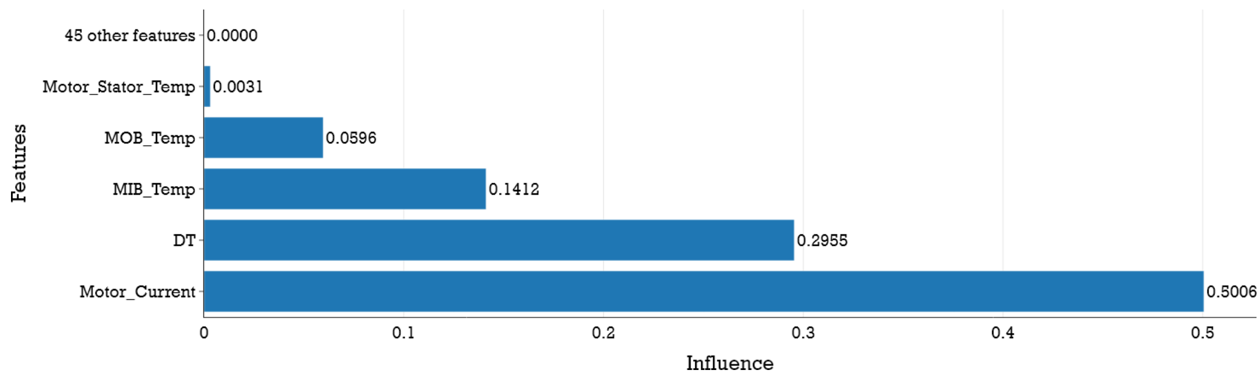

Figure 15 - Feature Influence on Predicting *misalignment*

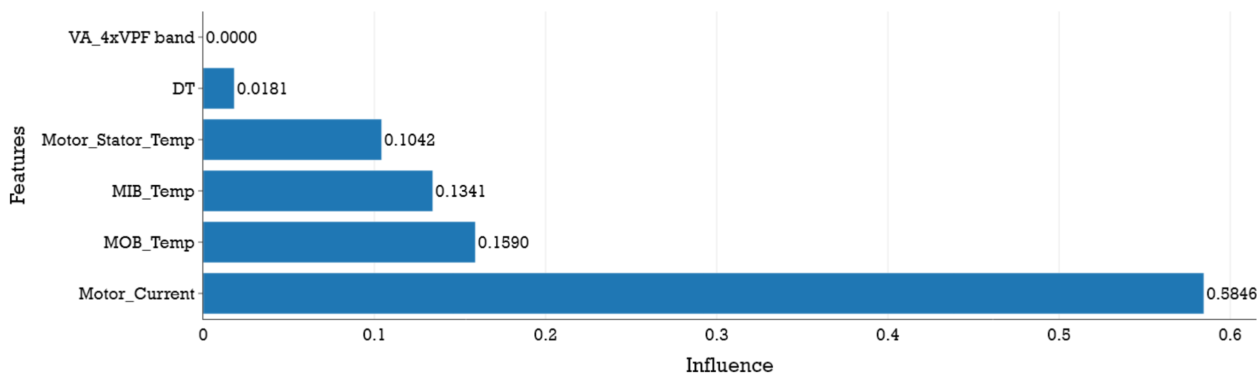

Figure 16 - Feature Influence on Predicting *waterbox fouling*

Essentially, when an instance of *unhealthy* pattern is fed to multiclass prediction model, the influence of features in assessing (i.e., estimating) the fault can be obtained from the SHAP values, and eventually the most significant features will be retained for the prognosis of an identified fault. For instance, the *CWP diffuser* fault was identified between April 01, 2020 and April 21, 2020. An instance, x from 10:00 AM on April 13, 2020, was fed into the fault prediction model, then its respective SHAP values for each fault were predicted, as shown in Figure 17 to Figure 21. Figure 17 to Figure 21, in which $E[f(x)]$ and $f(x)$ are the average prediction value and the output (i.e., logit) value for a fault, respectively. The term $f(x)$ for a fault is the linear combination of all associated SHAP values from each feature, plus the associated $E[f(x)]$ given by:

$$f(x) = E[f(x)] + \sum_{i=1}^{m} SHAP_i \tag{1}$$

where m is the total number of features used in the multiclass classifier input data.

Hence, in Figure 17 to Figure 21, the value inside the bars represents the SHAP values associated with the feature. The bar color red (blue) indicates the $+ve$ $(-ve)$ impact of the feature on $f(x)$. For a given fault scenario, the identified fault type will have the maximum $\mathrm{f}(x)$ compared to the other considered faults. For the considered instance, the $f(x)$ is maximum for *CWP diffuser* fault with the value of 2.74 (Figure 19), and all the input features contributed positively, with *MIB_1x_band* being the most dominant. After identifying the most dominant feature, its associated time-domain feature will be used by the prognosis model to forecast the trend of the identified fault. For the *CWP diffuser* fault, $\sigma_{VibXY-MIB}$ associated with *MIB_1x_band* will be considered in order to forecast the fault trend. The detailed discussion on the prognosis is discussed in the next section.

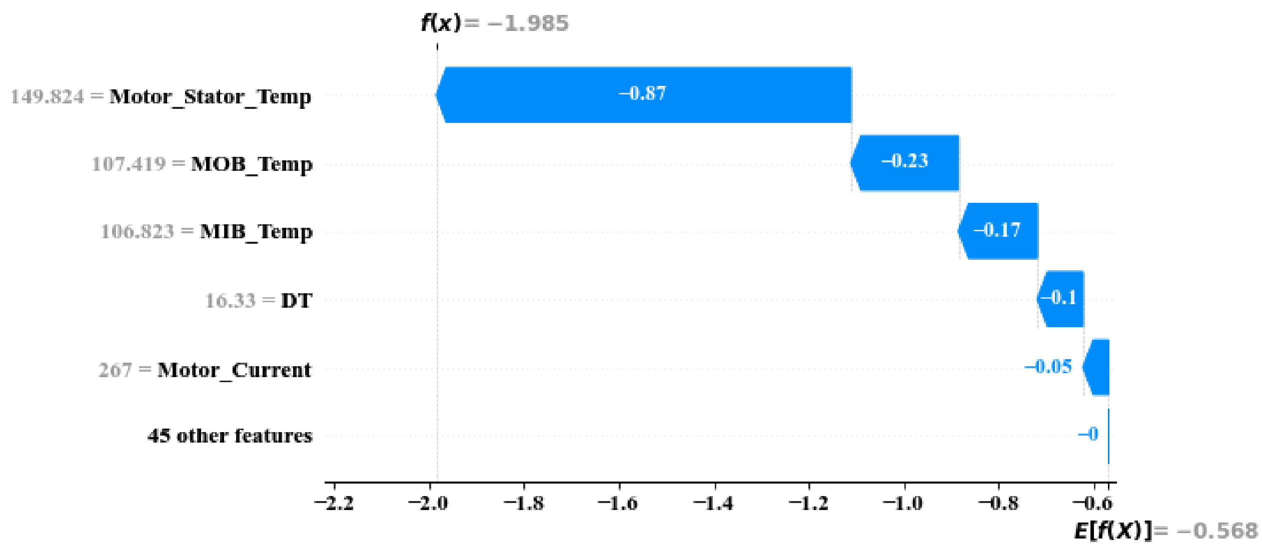

Figure 17 - SHAP Values for *air intake* for the Sample Associated with the *CWP diffuser* Fault

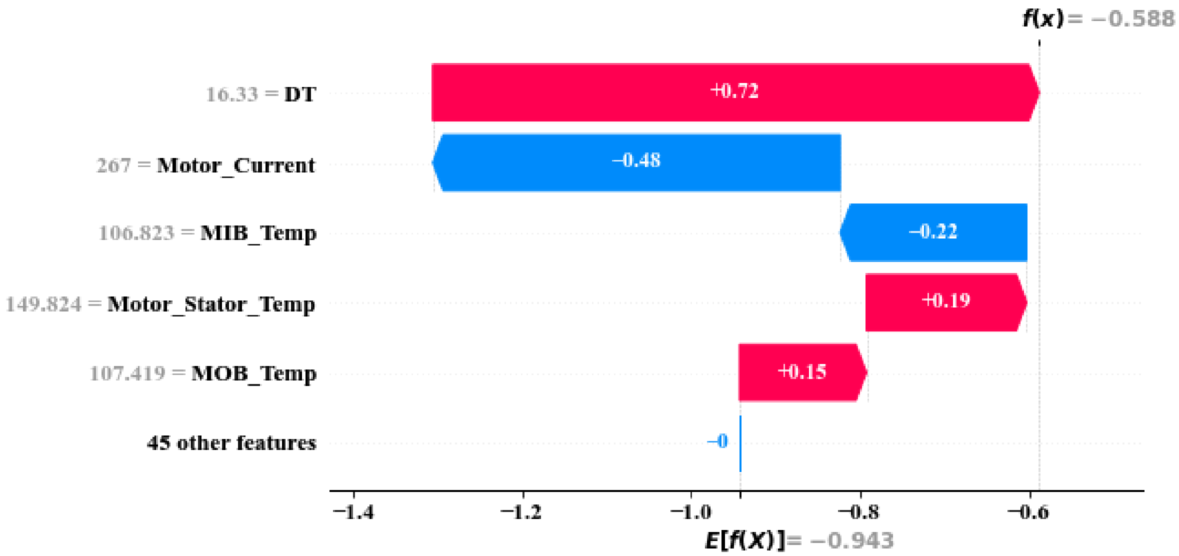

Figure 18 - SHAP Values for *bellmouth fail* for the Sample Associated with the *CWP diffuser* Fault

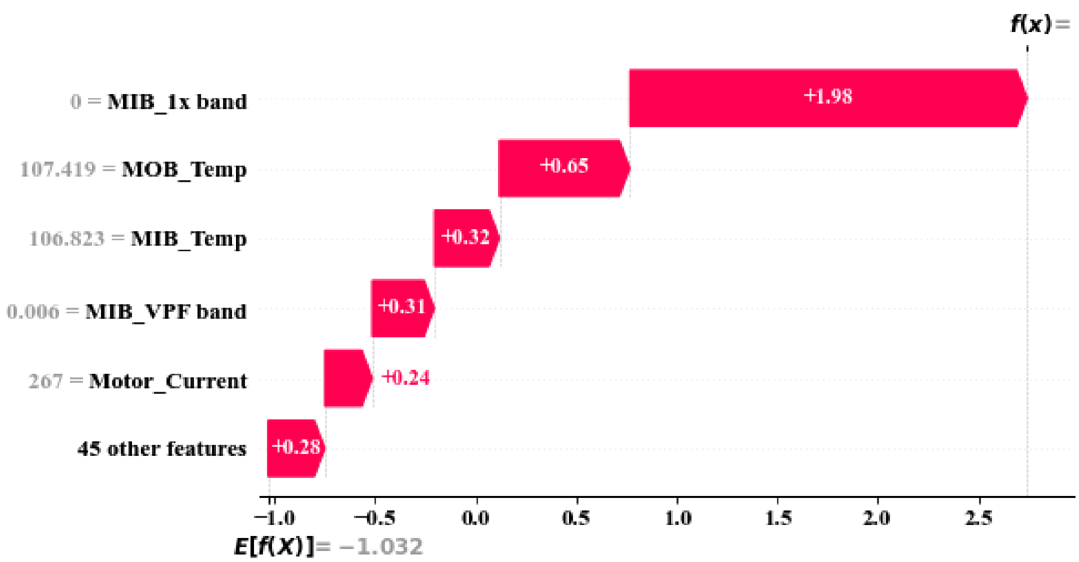

Figure 19 - SHAP Values for *CWP diffuser* for the Sample Associated with the *CWP diffuser* Fault

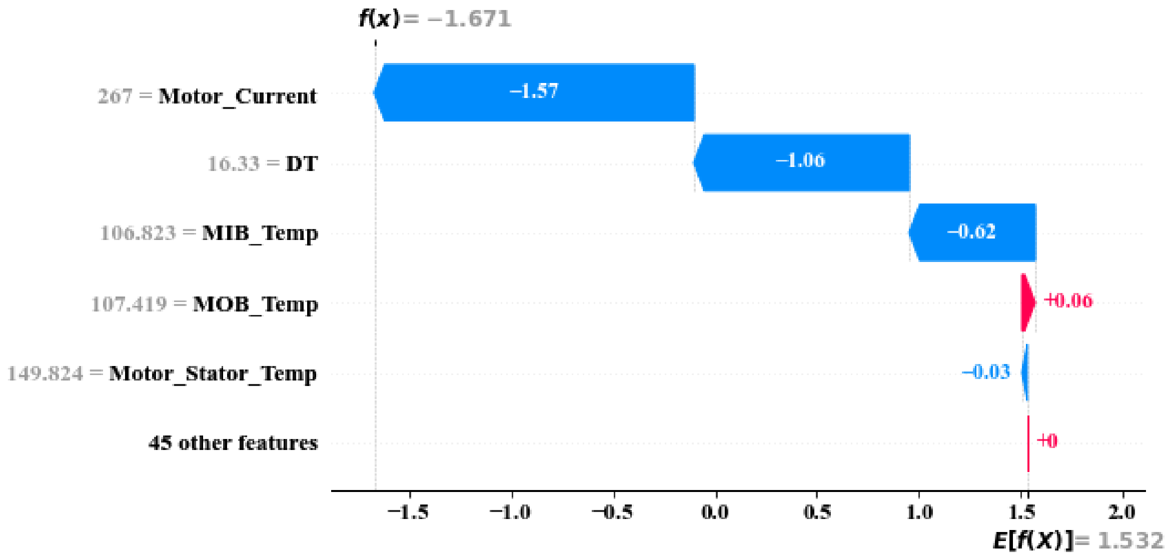

Figure 20 - SHAP Values for *misalignment* for the Sample Associated with the *CWP diffuser* Fault

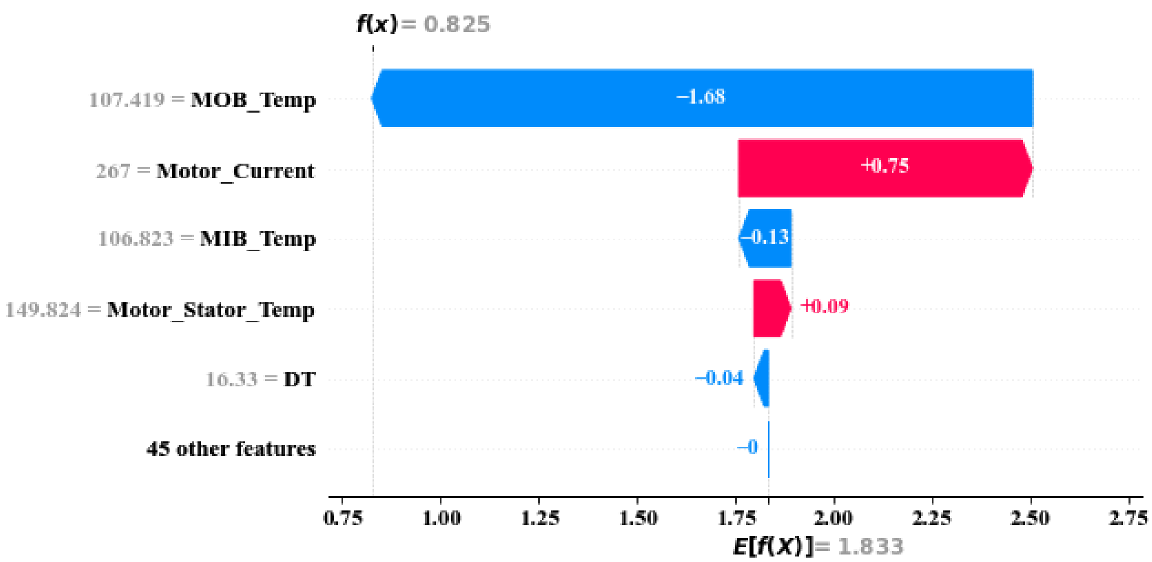

Figure 21 - SHAP Values for *waterbox fouling* for the Sample Associated with the *CWP diffuser* Fault

4.4 PREDICTIVE MODELING

The predictive modeling is part of the prognostic modeling effort. In this research, specifically, the focus is on short-term forecasting of developing degradations or faults. A maximum forecast of two days ahead is presented in this section. In some cases, long-term predictions are desired, but caution must be practiced because the uncertainty bounds around the forecasted value are expected to be large. Short-term forecasting is appropriate for this research because with the onset of a degradation (fault), it is important to track and address the fault before the condition causes components to reach their trip limits. Based on the historical data analysis, it is inferred that the time interval between the onset of a fault and it being addressed is at most four weeks (in most cases within a week). This predictive modeling approach does not support WO generation with PM planning, which is scheduled thirteen weeks in advance.

4.4.1 Short-term Forecasting

When forecasting time-series data, the aim is to estimate – as accurately as possible, considering all historical information and model parameters that could impact the forecast – how the sequence of observations will progress into the future [9]. In an industrial setting such as the power industry, predicting trends over the course of an hour to a week is desired, using short-term forecasting for cost-effective risk management. One such popular algorithm for short-term forecasting is auto-regressive integrated moving average (ARIMA).

A time series is stationary if its statistical properties (e.g., mean and variance) are constant over time. In the short term, the patterns always look identical in a statistical sense, and can be viewed as a linear combination of signal plus noise. The signal can reflect a pattern of fast or slow mean variation, with or without sinusoidal oscillation. It can also have a seasonal component. The ARIMA forecasting for a stationary time series is linear (i.e., regression type) with the predictor variable consisting of lags of the dependent variable and/or lags of the forecast errors, as given by:

$$\boldsymbol{predicted\ value, y_t = \mu + \epsilon_t + \underbrace{\textstyle\sum_{i=1}^{p} \gamma_i y_{t-i}}_{Auto-regressive\ (AR)} + \underbrace{\textstyle\sum_{j=1}^{q} \theta_j \epsilon_{t-j}}_{Moving\ Average\ (MA)}} \tag{2}$$

where μ is a constant, ϵ_t is the noise term, γ_i is the coefficient for the lagged variable at (*t-i*), and θ_j is the coefficient for the lagged error term in time (*t-j*).

In ARIMA, the lags of the stationarized series in Equation (2) are auto-regressive (AR) terms, lags of forecast errors are moving average (MA) terms, and the time-series that is differenced to be made stationary is said to be an integrated version. The ARIMA model is classified as an *ARIMA(p,d,q)* model, where *p* is the number of AR terms, *d* is the number of differences needed for becoming stationary, and *q* is the number of MA terms. The types of ARIMA model that are encountered are given in Table 8.

ARIMA(p,d,q)	Model type
***ARIMA*(1,0,0)**	**First-order AR model**
***ARIMA*(0,1,0)**	**Random walk**
***ARIMA*(1,1,0)**	**Differenced first-order AR model**
***ARIMA* (0,1,1) without constant, μ**	**Simple exponential smoothing**
***ARIMA* (0,1,1) with constant, μ**	**Simple exponential smoothing with growth**
***ARIMA* (0,2,1) or (0,2,2) without constant, μ**	**Linear exponential smoothing**
***ARIMA* (1,1,2) with constant, μ**	**Damped-trend linear exponential smoothing**

Table 8 - ARIMA Model Types

4.4.2 Walk-Forward Prognostic Model with ARIMA

Conventional methods of validation and cross validation are infeasible for time-series predictions. In particular, when the CWP undergoes a particular degradation, the trend of the selected feature will continue to change; hence, the ARIMA model can be trapped into under- or overfitting. Thus, one possible solution is a walk-forward approach [9] that retrains the model periodically or at a customized interval, incorporating newly collected samples. In doing so, data from too far in the past are excluded using a rolling window, as shown in Figure 22. At time instance $t = \{1,2,3,..\}$, the data collected for the last *D* days will be used to train ARIMA model and update its parameters *(p,d,q)* = (p_t, d_t, q_t). After instance t, the updated ARIMA model will be used for forecasting. Thus, at a customized interval, the ARIMA model is retrained and its model parameters updated, then it is used for forecasting.

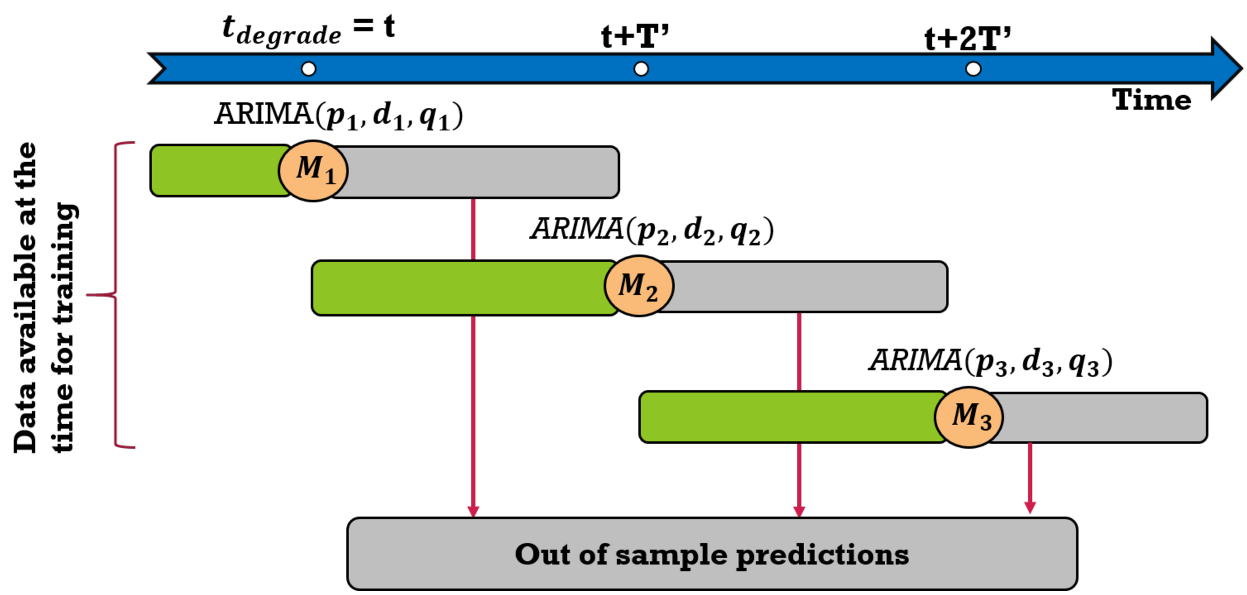

Figure 22 - Schematic Representation of a Walk-Forward Prognostic Model Using ARIMA

Thus, when an *unhealthy* condition is detected for CWP-{ij} at time instance t, the most dominant feature $F_{Features}(1,t)_{ij}$ from the fault prediction is extracted, and forecasting of that dominant feature is performed. At time instance t, the ARIMA model is trained using *healthy* data collected up to (*t-D*) hours, and these are used for out-of-sample predictions until τ hours, as shown in Figure 23. After every T' hours, the model is retrained using past *D* hour data, and the process repeats.

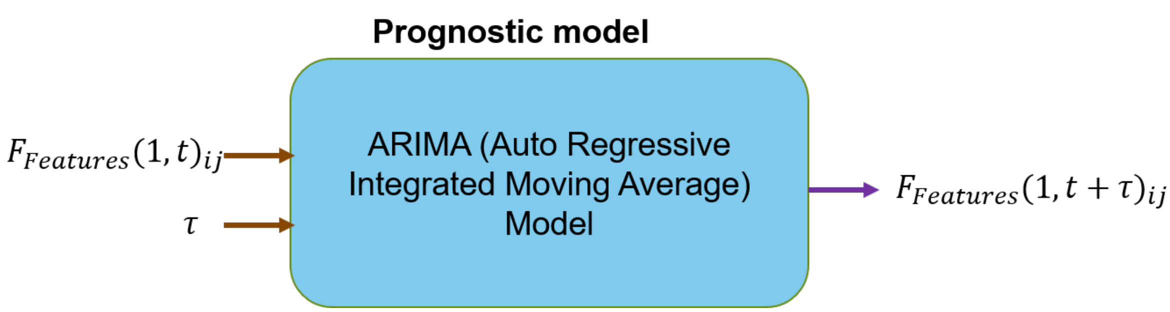

Figure 23 - Forecasting Using the ARIMA Model

4.4.3 Results and Discussion

For *CWP diffuser* case, the fault estimation model with SHAP values identified *MIB_1x_band* as the most significant feature. For prognosis, instead of forecasting the frequency-domain feature, its time-domain counterpart $\sigma_{VibXY-MIB}$ was considered, as it comprises the total magnitude of the vibration signal. For demonstration purposes, the *healthy* samples from April 29 to May 15 of 2020 were appended back between March 15 and March 31 of 2020, just before the *CWP diffuser* issue began on April 1, 2020, as shown in Figure 24.

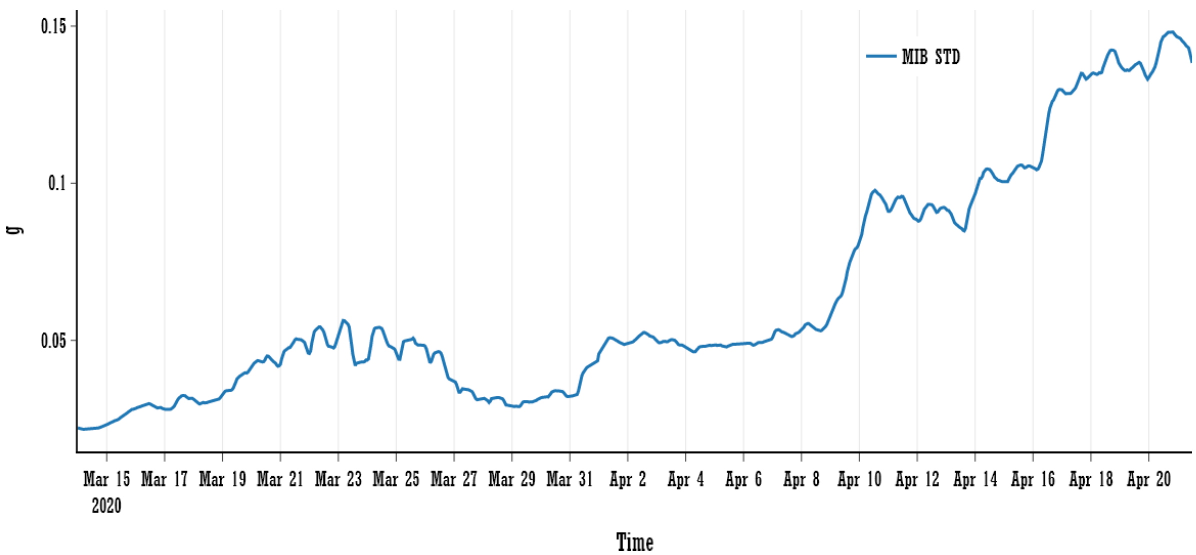

Figure 24 - $\sigma_{VibXY-MIB}$ Signal for Prognosis

The *CWP diffuser* issue was identified on April 1, 2020; hence, the previous seven days (which is optional) of data corresponding to the *healthy* condition were initially used to train the ARIMA model as shown in Figure 26. Similarly, while forecasting, the past seven days of collected samples were used each time to retrain the ARIMA model to update its parameters, *(p,d,q)*. The retraining is essential because the trends can change (see Figure 24) significantly, and thus the order parameters (*p,d,q*) will also change. An instance of ARIMA model retraining during *CWP diffuser* fault is shown in Figure 26.

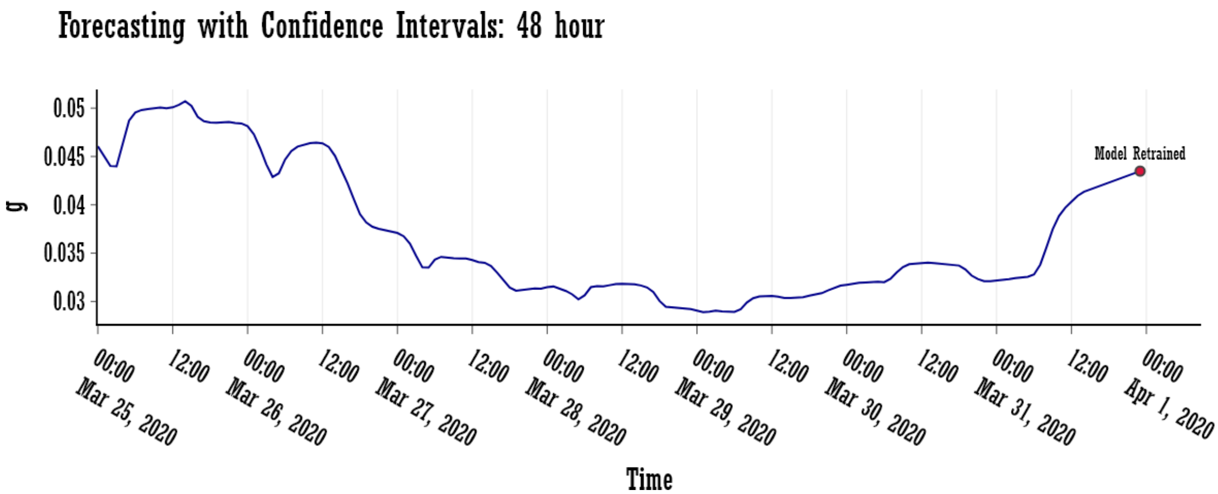

Figure 25 - Retraining Instance During Forecasting for March 31, 2020

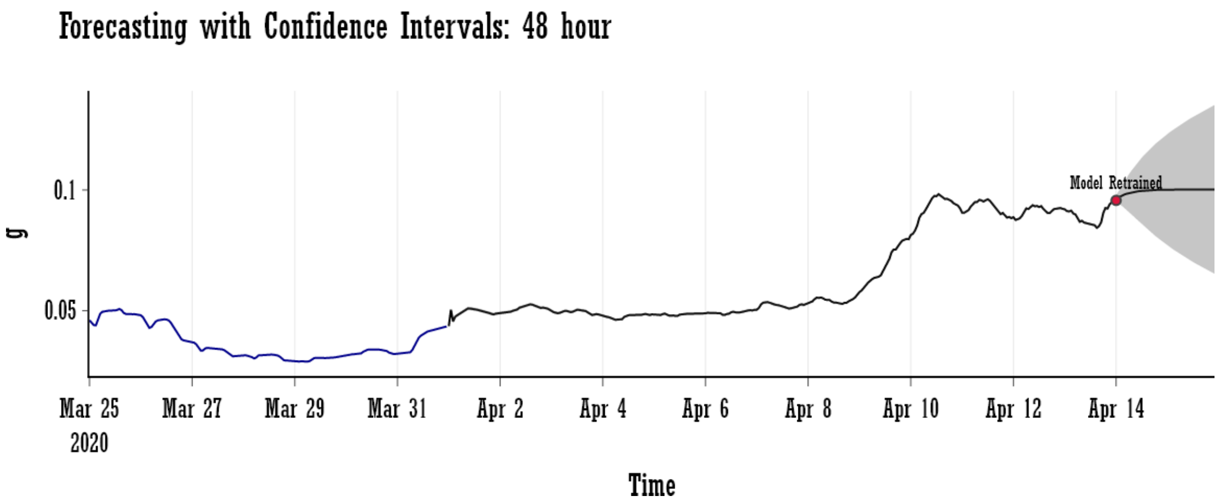

Figure 26 – Retraining Instance During Forecasting for April 14, 2020

At each time instance, the dominant parameter is forecasted for 48 hours (can be set as per business needs). The mean value of the forecasted trends for $\sigma_{VibXY-MIB}$ for *CWP diffuser* issue are shown in Figure 27. As the forecasting steps increase (here, up to 48 hour) the uncertainties in the forecasted value increase with wider confidence intervals (the grey regions in Figure 26). For the longer forecasting steps, the spikes in prediction are due to the length of the past data considered in the model re-training, as well as the amount of fault trends captured. From Figure 27, it is evident that until April 9, 2020, the fault trend had a slow variation, but shot up starting on April 9, 2020. Hence, the retrained ARIMA models after April 10, 2020 have fewer residuals, even for longer forecasting steps. This can be further verified by plotting each forecast trends with its 95% confidence interval, as shown in Figure 31 to Figure 31.

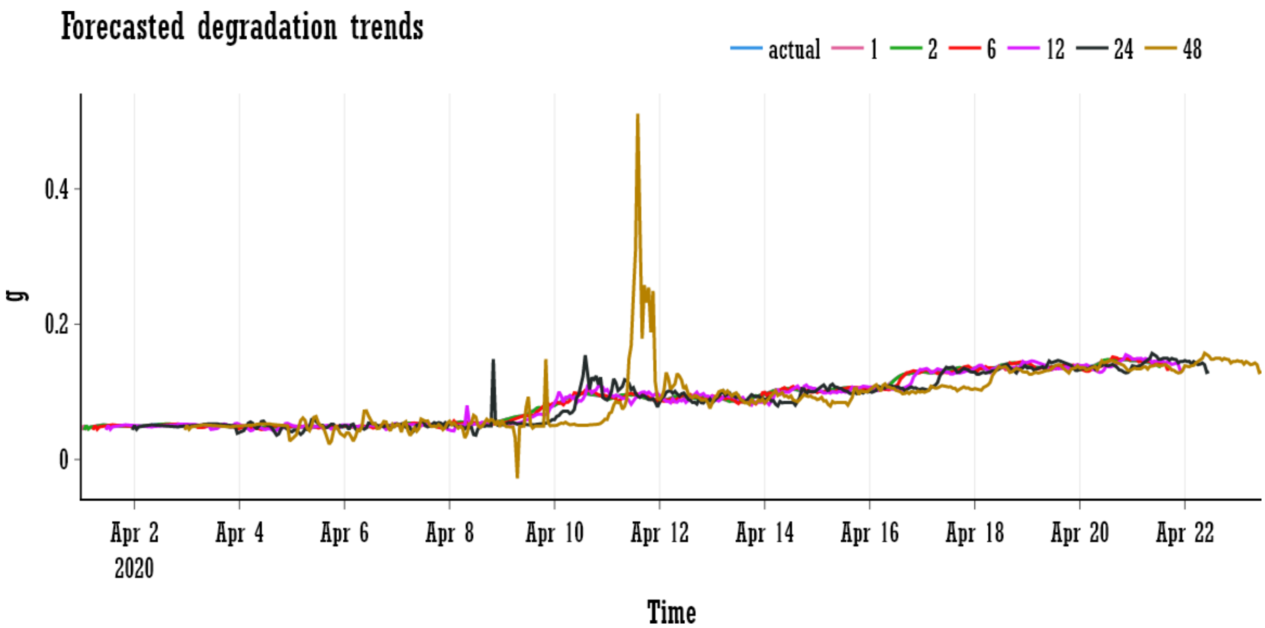

Figure 27 - Forecast Trends for Feature $\sigma_{VibXY-MIB}$ for 1st, 2nd, 6th, 12th, 24th, and 48th hours

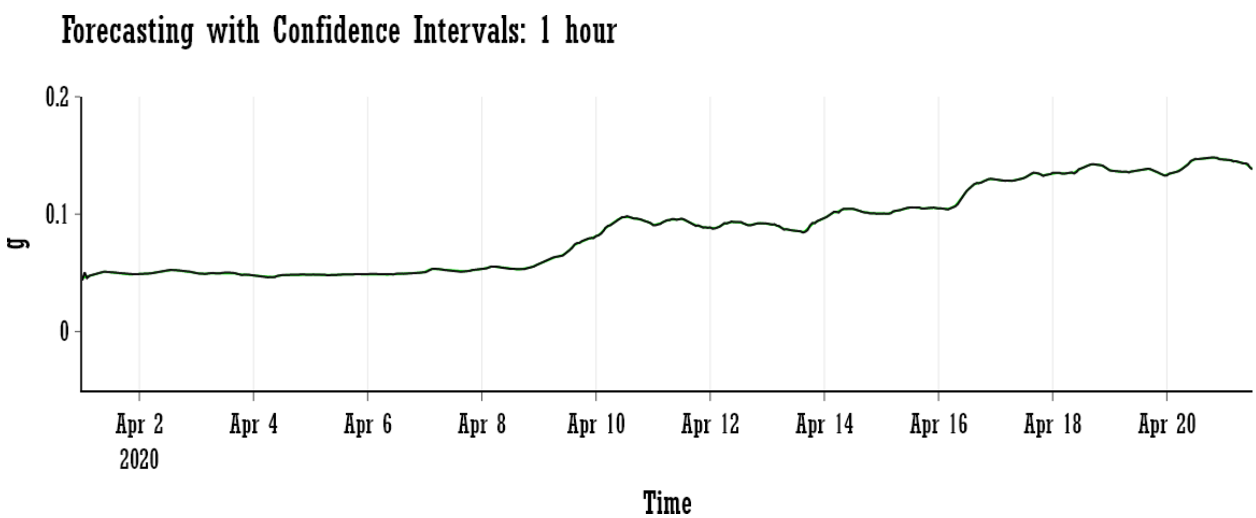

Figure 28 - Forecasting Trends with a 95% Confidence Interval for the 1st Hour

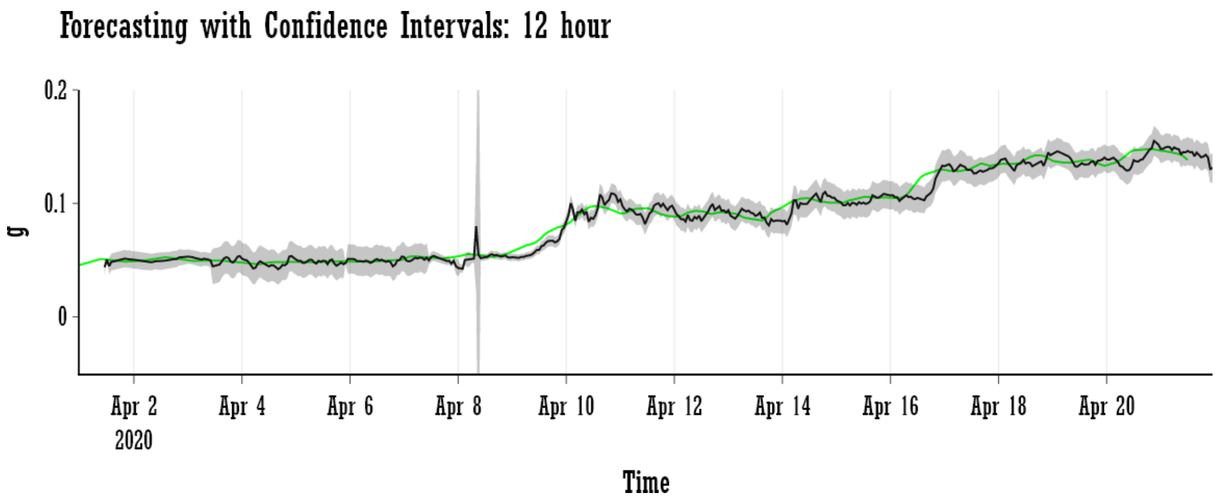

Figure 29 - Forecasting Trends with a 95% Confidence Interval for the 12th Hour

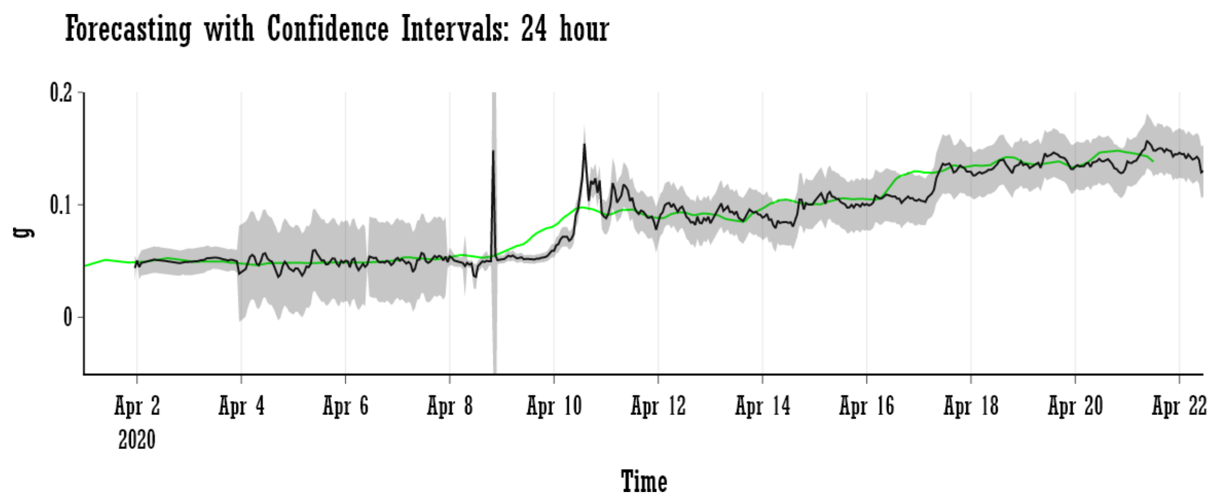

Figure 30 - Forecasting Trends with a 95% Confidence Interval for the 24th Hour

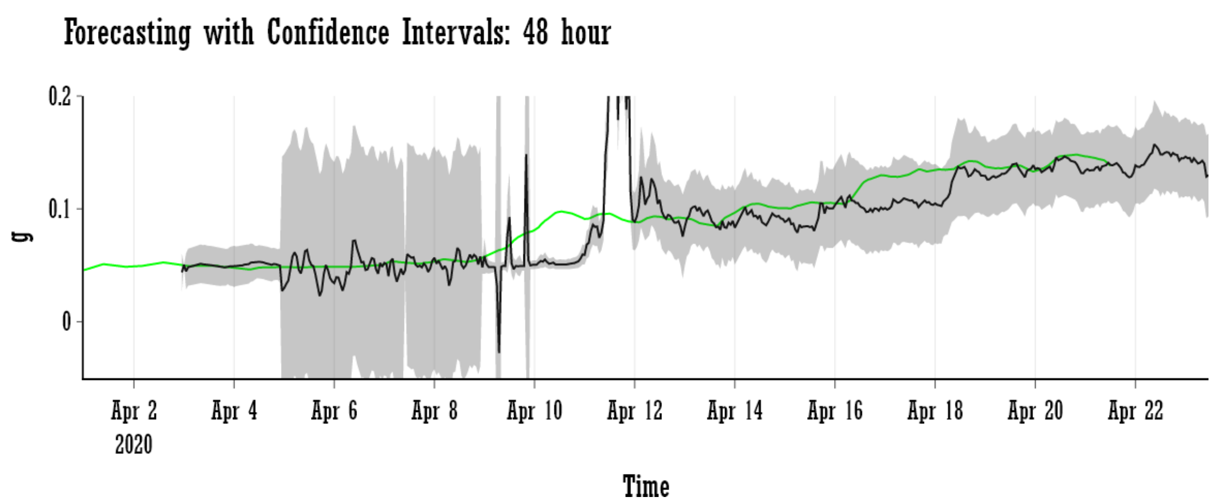

Figure 31 - Forecasting Trends with a 95% Confidence Interval for the 48th Hour

In addition, the trained model performance was analyzed using residuals (defined as the difference between observations and the corresponding fitted values). The residuals are useful for verifying whether the ARIMA model adequately captured the information within the data. For good forecasting, the residuals should be uncorrelated, and have zero mean, otherwise the forecasts will be biased. The residual plot for the ARIMA model for forecasting the parameter $\sigma_{VibXY-MIB}$ is shown in Figure 32. The residual seems to fluctuate around zero and have a uniform variance (Figure 32, top left), while the density plot (Figure 32, top right) suggests residual is normally distributed with a zero mean. Since a few dots do not align with the red line, the residual distribution is skewed (Figure 32, bottom left). Finally, the correlation plot indicate that the residuals are not uncorrelated (Figure 32, bottom right).

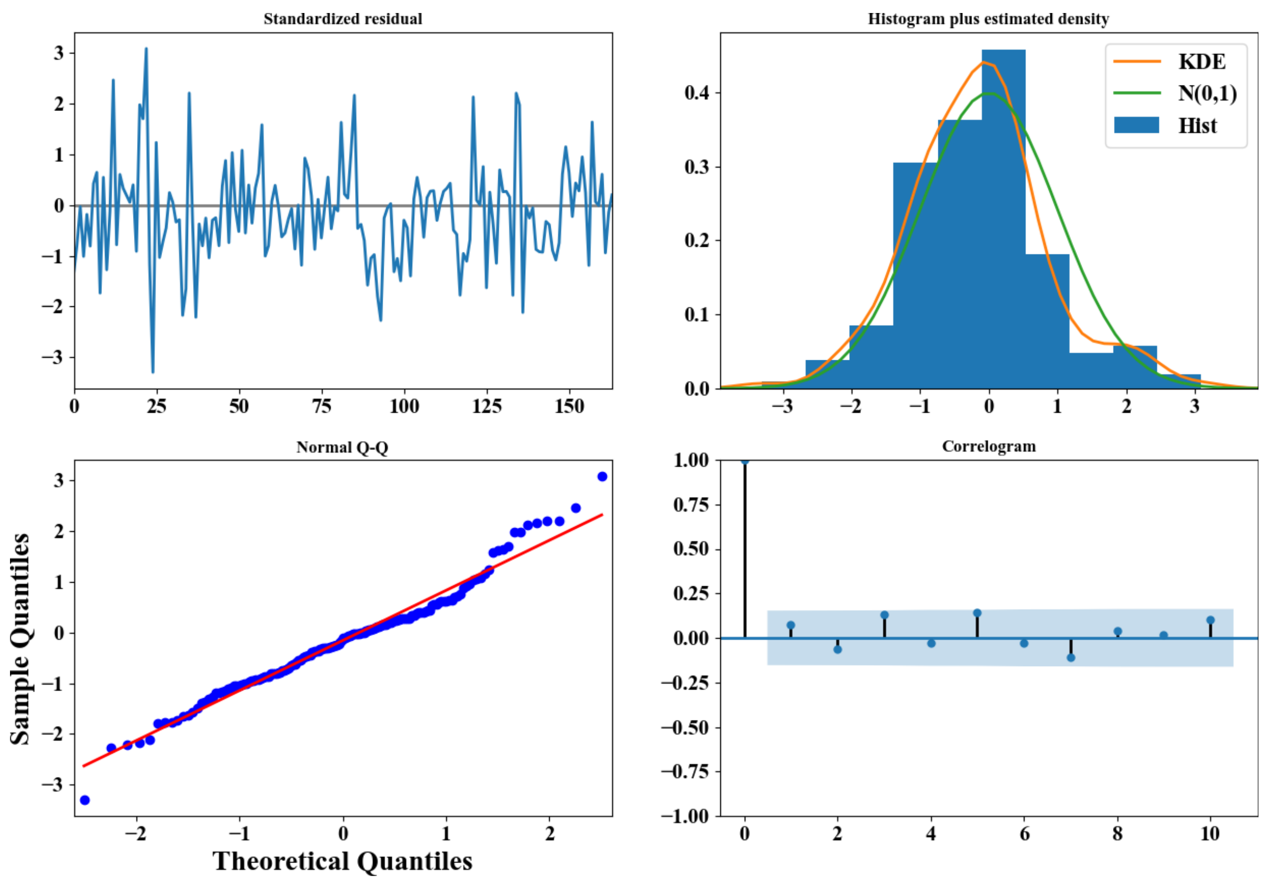

Figure 32 - ARIMA Residual Chart.

4.4.4 Calculation of Degradation Index

From the forecasted values, the level of degradation in terms of degradation index, β is calculated by:

$$\beta = 1 + \frac{\left|F_{Features}(1, t + \tau)_{ij} - E\left[F_{Features}(1, t - \delta)_{ij}\right]\right|}{E\left[F_{Features}(1, t - \delta)_{ij}\right]} \tag{3}$$

where $F_{Features}(1, t + \tau)_{ij}$ is the forecasted value of $F_{Features}(1, t)_{ij}$ at hour $t + \tau$, $E\left[F_{Features}(1, t - \delta)_{ij}\right]$ is the average value of the feature $F_{Features}(1, t)_{ij}$ when CWP-{ij} is in *healthy* condition, and δ is time duration prior to degradation. Thus, the estimated β parameter will be used by hazard function for risk modeling and economic formulation, as discussed in detail in the coming sections.

4.5 HAZARD AND ECONOMIC MODELING

This section describes the application of the continuous-time, discrete-states Markov chain model to the modeling of CWS motor and pump's [12] CM and PM tasks. Some of the information is this section follows previously published information presented in [13]; however, it is reiterated for the convenience of the reader. For modeling purposes, several assumptions were made about the M&P sets and their operation. Specifically, it was postulated that each M&P set could be in one of the three states shown in Figure 33. The set can be fully operational and running, or may be undergoing CM or PM. Transitions between states are governed by transition rates – namely, the following model parameters: failure/downtime rate $\lambda(t)$, CM rate $\mu(t)$, PM scheduling rate $\eta(t)$, and PM rate $\nu(t)$. If the rates are not time-dependent, the model is called a "homogeneous Markov chain." For time-dependent rates, the model becomes a "non-homogeneous Markov chain." Transitions between states may happen at random times, and we are interested in calculating the probabilities of different states, given the transition rates.

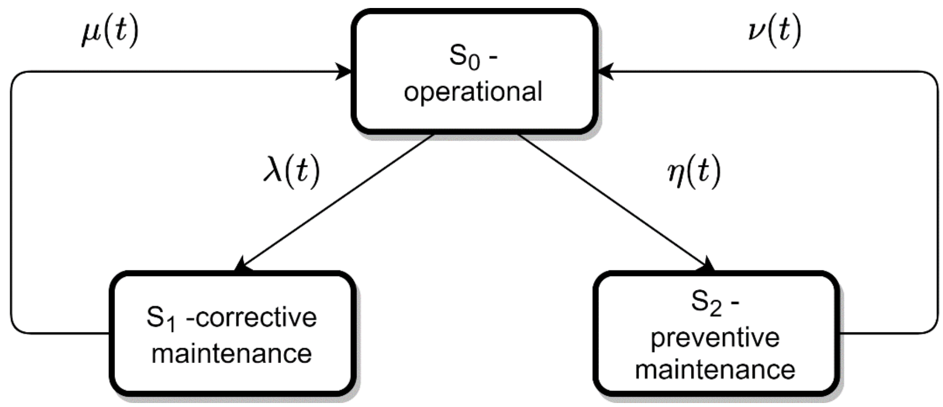

Figure 33 - Transition Diagram for a Three-State Model of a Single Motor-Pump Set

Under these conditions, the nonstationary or nonhomogeneous Markov chain with discrete states and continuous time can be described by a system of Chapman-Kolmogorov equations [14]:

$$\frac{dP_i(i)}{dt} = \sum_{j=1}^{N} P_j(t) \cdot \lambda_{j,i}(t) - P_i(t) \cdot \sum_{j=1}^{N} \lambda_{ij}(t) \tag{4}$$

with initial and normalization conditions $p_0(0) = 1, \sum_{i=1}^{N} P_i\,(t) = 1 \;\; for\ any\ t.$

The system of differential equations that, governing the evolution of the three-state model shown in Figure 33 can be written as follows [14]:

$$\frac{dp_0(t)}{dt} = \mu(t) \cdot p_1(\mathrm{t}) - \lambda(t) \cdot p_0(t) + \nu(t) \cdot p_2(\mathrm{t}) - \eta(t) \cdot p_0(t)$$

$$\frac{dp_1(t)}{dt} = \lambda(t) \cdot p_0(\mathrm{t}) - \mu(t) \cdot p_1(t) \tag{5}$$

$$\frac{dp_2(t)}{dt} = \eta(t) \cdot p_0(t) - \nu(t) \cdot p_2(t)$$

with the following initial and normalization conditions: $p_0(0) = 1, p_0(t) + p_1(t) + p_2(t) = 1.$

If the three-state model is homogeneous with time-independent rates, as shown in Figure 34, the final probabilities of the three states can be calculated analytically as follows [15]:

$$p_1 = \frac{1}{1+\frac{\mu}{\lambda}+\frac{\mu\cdot\eta}{\lambda\cdot\nu}} ; p_0 = \frac{\mu}{\lambda} p_1 ; p_2 = \frac{\eta}{\nu} p_1 \quad (6)$$

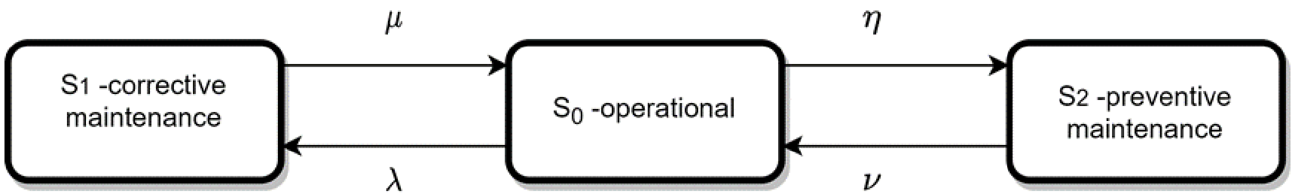

Figure 34 - Transition Diagram for a Three-State Homogeneous Model of a Single Motor-Pump Set.

However, for the non-homogeneous model, the system of equations must be solved numerically. Either way requires availability of the model's parameters (i.e., transition rates), which can be estimated from the plant's historical operational data as described in [1].

For a homogeneous model with constant failure rates, the probability distribution of the time intervals is assumed to be exponential, with a probability density function (PDF) – f(t), cumulative density function (CDF) – F(t), and reliability functions R(t), as shown in Figure 35 for λ=1.0·10-3. The exponential PDF can be written as:

$$f(t) = \lambda e^{-\lambda t}, t > 0 \quad (7)$$

where λ is the parameter of the distribution. Assuming an exponential distribution for the times between events, the probability of an event prior to time T, as well as an event between times t_1 and t_2 can be written as follows [14]:

$$P(T \leq t) = \int_0^t \lambda \cdot e^{-\lambda \cdot s} ds = 1 - e^{-\lambda \cdot t} = F(t); P(t_1 \leq T \leq t_2) = \int_{t_1}^{t_2} \lambda \cdot e^{-\lambda \cdot s} ds = e^{-\lambda \cdot t_1} - e^{-\lambda \cdot t_2} \quad (8)$$

The reliability function [16] for the exponential distribution is defined as:

$$R(t) = P(T > t) = 1 - F(t) = 1 - \left(1 - e^{-\lambda \cdot t}\right) = e^{-\lambda \cdot t},\ F(t) + R(t) = 1 \quad (9)$$

The reliability function R(t) is a complementary value to the CDF. While CDF is the probability that an M&P set will fail (or degrade) by time t, R(t) is the probability that an M&P set will survive at least up to time t. Any time scale can be used for exponential distribution and this is the reason for not specifying time units in the subsequent figures. The specific parameter values used in the description of Markov model and its distributional assumptions are used in this report for demonstration of underlying principles and in practice should be replaced with parameters reflecting the current state of a CWS for a specific unit.

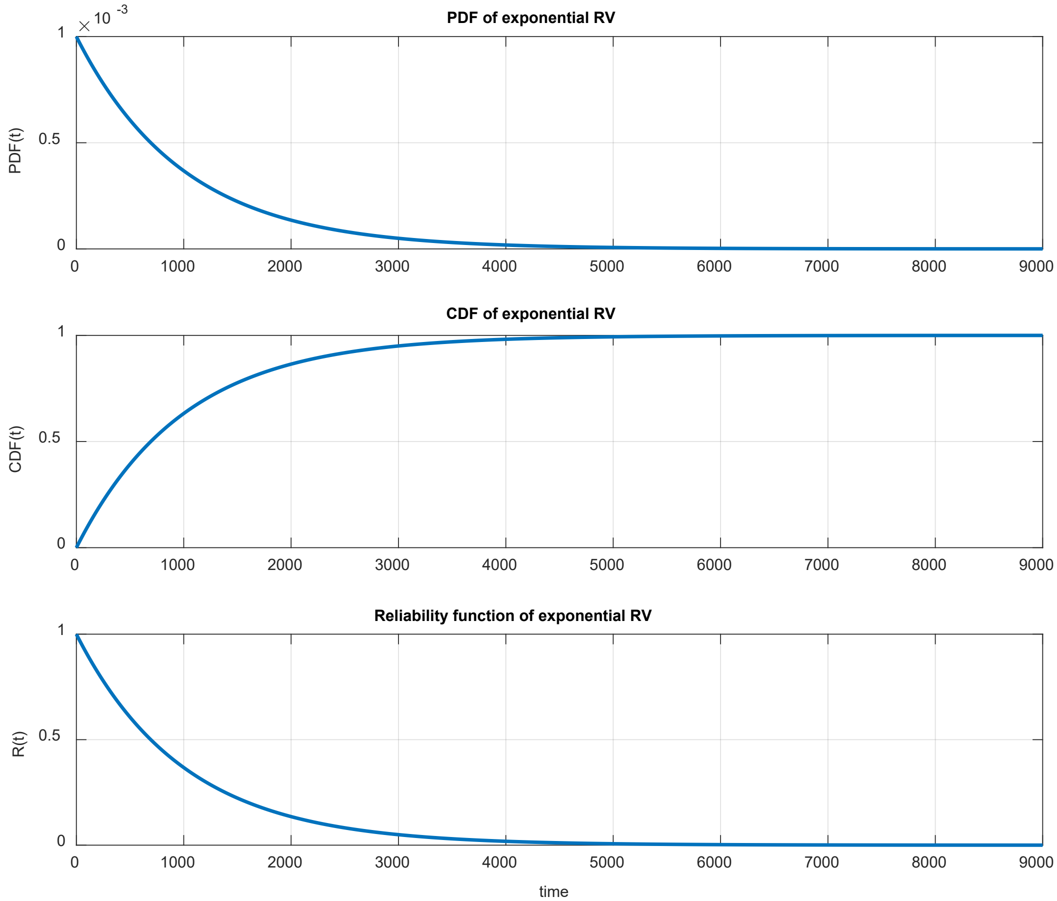

Figure 35 - Probabilistic Functions for Exponential Random Variable (RV)

Assuming an exponential distribution, the instantaneous unconditional probability of failure for a brand new (or as good as new) M&P set within time interval dt after starting can be calculated as [16]:

$$\boldsymbol{F(t) = P(0 \leq T \leq 0 + dt) = F(t) = \int_{0}^{0+dt} \lambda \cdot e^{-\lambda \cdot s} ds\,, = 1 - e^{-\lambda \cdot dt}}\,;$$

$$\boldsymbol{F(t) = P(0 \leq T \leq 0 + dt | T \geq 0) = \frac{P(0 \leq T \leq 0 + dt\,\&T \geq 0)}{P(T \geq 0)}}$$

$$\boldsymbol{= \frac{P(0 \leq T \leq 0 + dt)}{P(T \geq 0)} = \frac{1 - e^{-\lambda \cdot dt}}{1} = 1 - e^{-\lambda \cdot dt}} \qquad \textbf{(10)}$$

As seen, probability of the M&P set failing within time interval dt is $1 - e^{-\lambda \cdot dt}$. This can be called the "unconditional probability of failure" as we assumed that the set was started at time stamp zero. In practice, however, we are often interested in the probability of failure conditional on some run-time interval. In other words, we are interested in the probability of failure of the M&P set after it survived some time interval t. In this case, the conditional probability of the failure can be calculated as [16]:

$$F(t) = P(t \leq T \leq t + dt | T \geq t) = \frac{P(t \leq T \leq t + dt \,\&T \geq t)}{P(T \geq t)} = \frac{P(t \leq T \leq t + dt)}{P(T \geq t)}$$
$$= \frac{e^{-\lambda \cdot t}\left(1 - e^{-\lambda \cdot dt}\right)}{1 - F(t)} = \frac{e^{-\lambda \cdot t}\left(1 - e^{-\lambda \cdot dt}\right)}{e^{-\lambda \cdot t}} = 1 - e^{-\lambda \cdot dt} \tag{11}$$

As is seen from Equation (8), the conditional probability of failure after surviving time t is the same as the unconditional probability [Equation (7)] of failure right after starting. This is the memoryless property of exponential distribution: the probability that a brand-new motor will fail after dt hours is the same as that of failure within dt hours after running for t hours. In reliability research, the unconditional probability of failure is often called "failure to start", while the conditional probability of failure is called "failure to run". The memoryless property of exponential distribution also implies that the failure rate is unchanging over time. The conditional CDF can be written as [16]:

$$F_{T|T \geq t}(t) = P(t \leq T \leq t + dt | T \geq t) = \frac{P(t \leq T \leq t + dt \,\&T \geq t)}{P(T \geq t)} = \frac{P(t \leq T \leq t + dt)}{P(T \geq t)}$$
$$= \frac{F(t + dt) - F(t)}{P(T \geq t)} = \frac{e^{-\lambda \cdot t}\left(1 - e^{-\lambda \cdot dt}\right)}{1 - F(t)} = \frac{e^{-\lambda \cdot t}\left(1 - e^{-\lambda \cdot dt}\right)}{e^{-\lambda \cdot t}} = 1 - e^{-\lambda \cdot dt} \tag{12}$$

Differentiating the conditional CDF with respect to time t, we obtain the conditional PDF which can be written as:

$$\frac{F_{T|T \geq t}(t)}{dt} = \frac{P(t \leq T \leq t + dt | T \geq t)}{dt} = \frac{P(t \leq T \leq t + dt \,\&T \geq t)}{P(T \geq t)dt} = \frac{P(t \leq T \leq t + dt)}{P(T \geq t)dt}$$
$$= \frac{1}{P(T \geq t)} \frac{P(t \leq T \leq t + dt)}{dt} = \frac{f(t)}{P(T \geq t)} = \frac{\lambda \cdot e^{-\lambda \cdot t}}{e^{-\lambda \cdot t}} = \lambda,\ T \geq t \tag{13}$$

Equation (10) shows the failure rate to be independent of time, and is a constant equal to λ. This result demonstrates that the exponential distribution cannot be used to model aging or degradation, as it implies constant rates.

Unconditional and conditional CDFs, PDFs, and R for the exponential distribution and two different survival times are shown in Figure 36. As seen in the figure, the shapes of probabilistic functions do not change – the functions just shifted to the right of the survival time value. It can be shown analytically that the unconditional and conditional instantaneous failure rates are the same for the exponential distribution [16] and are equal to λ:

$$f(0 + dt) = \lambda \cdot e^{-\lambda \cdot dt}; f(0) = \lambda \cdot e^{-\lambda \cdot 0}\ for\ dt = 0$$
$$= \lambda - unconditional\ instanteneous\ failure\ rate\ for\ a\ brand - new\ motor$$

$$f_{T>200}(200 + dt) = \frac{\lambda \cdot e^{-\lambda \cdot (200+dt)}}{e^{-\lambda \cdot 200}} = \lambda \cdot e^{-\lambda \cdot dt}; for\ dt = 0; \lambda \cdot e^{-\lambda \cdot 0}$$
$$= \lambda - instanteneous\ failure\ rate\ conditioned\ on\ motor\ surviving\ 200\ hours$$

$$f_{T>500}(500 + dt) = \frac{\lambda \cdot e^{-\lambda \cdot (500+dt)}}{e^{-\lambda \cdot 500}} = \lambda \cdot e^{-\lambda \cdot dt}; for\ dt = 0; \lambda \cdot e^{-\lambda \cdot 0}$$
$$= \lambda - instanteneous\ failure\ rate\ conditioned\ on\ motor\ surviving\ 500\ hours$$

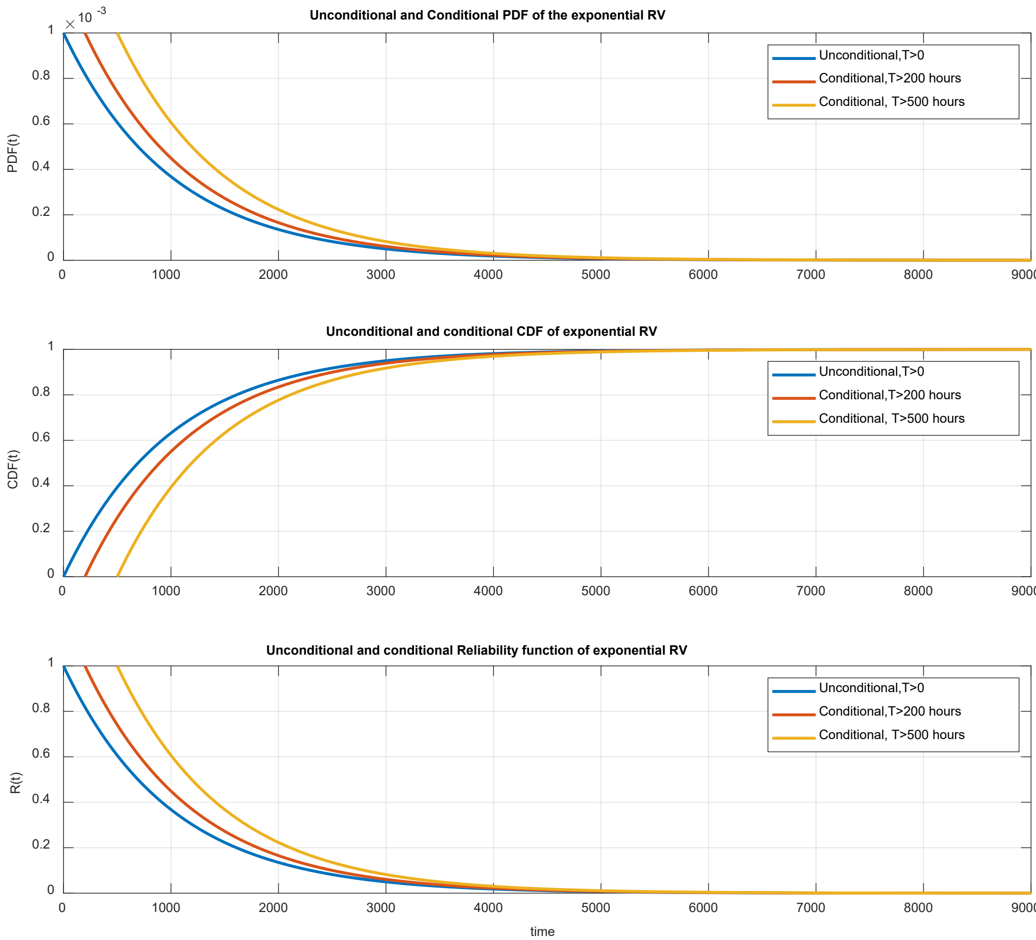

Figure 36 - Unconditional and Conditional PDFs, CDFs, and R for the Exponential Distribution

For the exponential distribution, the instantaneous failure rate is not time-dependent, but rather a constant. The exponential distribution cannot be used to model aging or degradation. The unconditional PDF is an unconditional predictor, and the conditional PDF is a conditional predictor. For the exponential distribution, they are identical. Knowing the past does not change our predictions if the exponential model is assumed.

The hazard function (or instantaneous failure rate) is a function of t, the time reached/survived without failure [16]**.** It is not a legitimate probability density, since its integral diverges. The instantaneous failure rate or hazard function is a conditional PDF calculated at survival time t.

Since the exponential distribution is inappropriate for modeling aging and degradation, other distributions must be explored to describe ongoing degradation. One of such widely used distributions is the Weibull distribution [16]:

$$\boldsymbol{f(t;\alpha,\beta) = \frac{\beta}{\alpha}\cdot\left(\frac{t}{\alpha}\right)^{\beta-1}\cdot e^{-\left(\frac{t}{\alpha}\right)^{\beta}};\alpha,\beta > 0, t \geq 0} \qquad \textbf{(14)}$$

The Weibull distribution has two parameters, α and β, and can be considered an extension of the exponential distribution [16]. Parameter β controls the shape of the distribution, while α controls the scale or variance as evident from Figure 37. The parameters of the Weibull distribution shown in Figure 37 are used to demonstrate the parametric dependency of the distribution and they are not obtained by fitting plant-specific data.

$$\boldsymbol{f(t;\alpha,\beta) = \frac{\beta}{\alpha}\cdot\left(\frac{t}{\alpha}\right)^{\beta-1}\cdot e^{-\left(\frac{t}{\alpha}\right)^{\beta}}\,;\alpha,\beta > 0, t \geq 0} \quad \textbf{(15)}$$

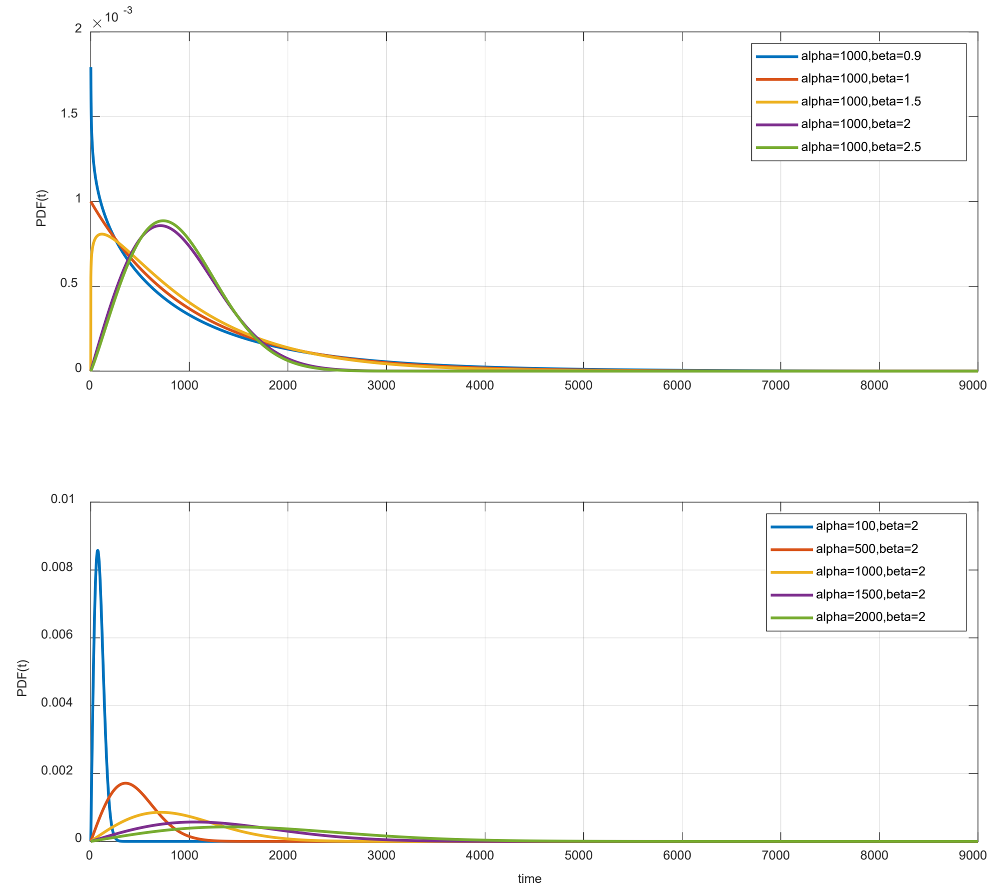

Figure 37 - Weibull PDF for Different Parameter Values of Parameters α and β

As seen from the top panel of Figure 37, if parameter α is fixed and β is varied, the PDF of the distribution is changing shape. If β is fixed and α is varied, as shown in the bottom panel, the PDF's shape remains the same, but the spread of the PDF changes, indicating changes in scale or variance. For the Weibull distribution, different values of β produce different distributions. For example, for β=1, the Weibull distribution becomes an exponential distribution, while for β=2, it corresponds to the Rayleigh distribution. The hazard function for the Weibull distribution can be written as [16]:

$$\boldsymbol{h(t) = \frac{\beta}{\alpha}\left(\frac{t}{\alpha}\right)^{\beta-1} - power\ of\ t} \quad \textbf{(16)}$$

The Weibull distribution includes three parametric models; one for the hazard function-power of t (increasing and decreasing), one linear (increasing-Rayleigh distribution), and one constant (exponential distribution) [16]. The Weibull hazard function for different values of parameter β is shown in Figure 38. The figure shows the hazard rate behavior for 900 hours; however, the behavior is retained if the time scale is extended to thousands of hours or days.

For example, for β=1, the hazard rate will be a constant, while for β=2, it will be a linear function of time regardless of the time scale.

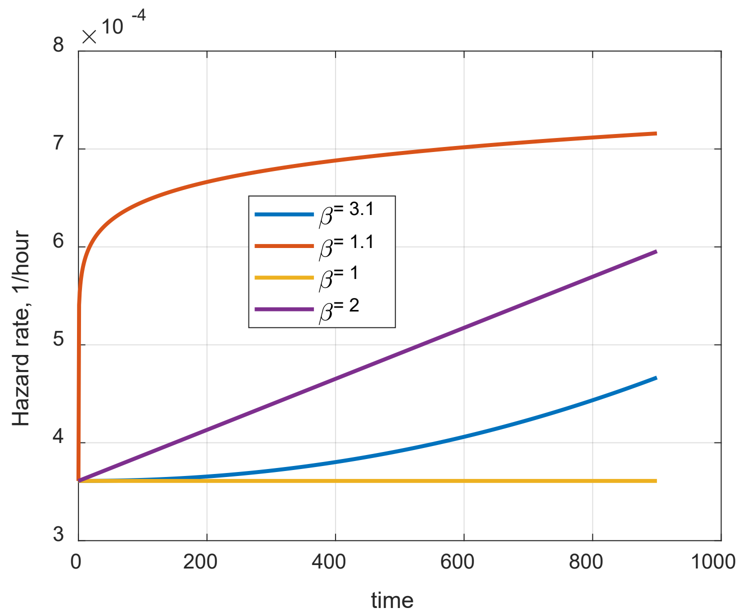

Figure 38 - Hazard Rate for the Weibull Distribution with Different Values of Parameter β and Parameter $\alpha = 2.7 \cdot 10^3$

In contrast to the exponential distribution, the Weibull distribution can be used to model aging and degradation, as its conditional PDF differs from its unconditional PDF, as shown in Figure 39. The figure is used to demonstrate the dependency regardless of the time scale.

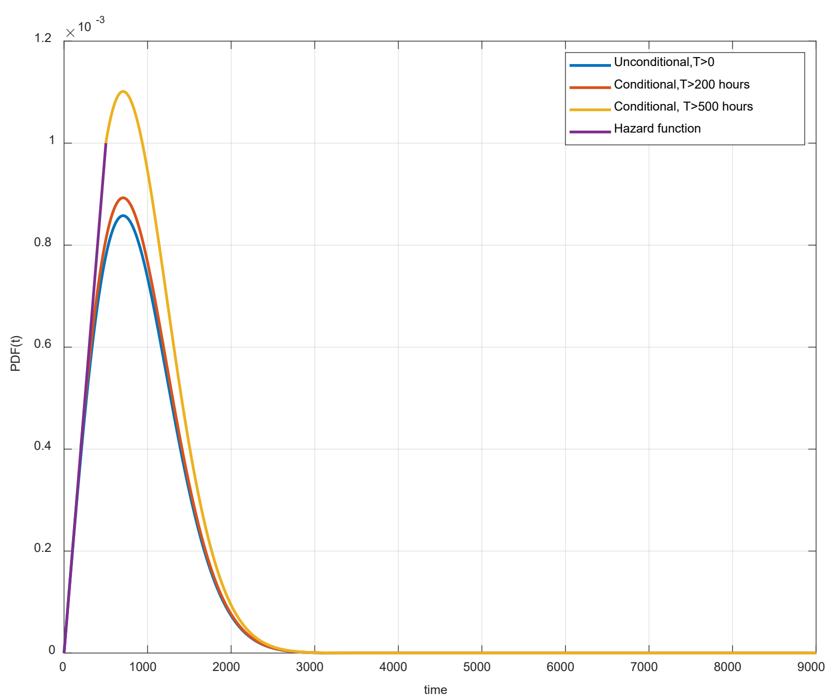

Figure 39 - Unconditional and Conditional PDFs and Hazard Function for Weibull Distribution with α=1000 and β=2

The Weibull distribution can also be used to model situations in which system reliability can improve over time, as shown in Figure 40. For parameter β values of less than one, the hazard function is decreasing. Figure 40 demonstrates that for β=0.9, the hazard rate is decreasing. Since the hazard rate is instantaneous failure rate, for Salem, this parameteriation can be used to describe the set behaviors after PM when the set can be regarded as "reliable as new".

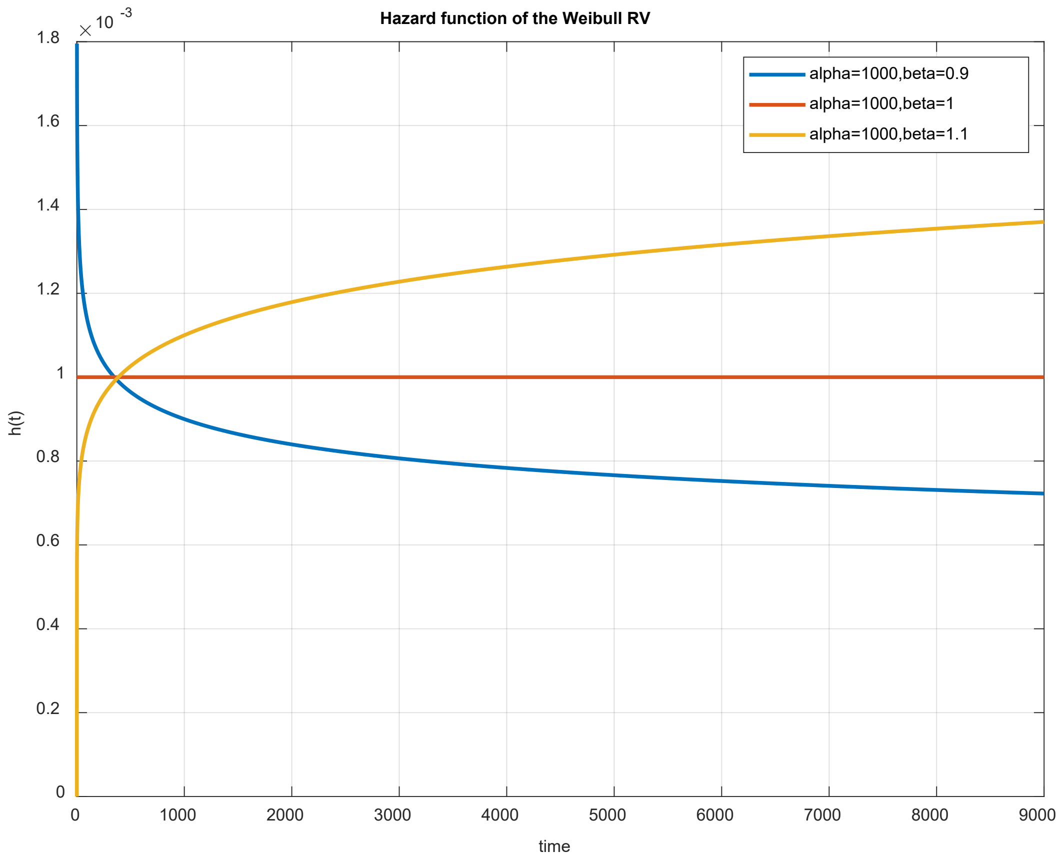

Figure 40 - Hazard Function of Weibull Distribution for Different Values of Parameter β and α=1000.

The Weibull model was been employed to model the degradation of an M&P set at the Salem NPP [2]. Figure 41 shows time-dependent rates for the case of M&P set degradation. The data plotted in the figure are real plant data and the top panel shows the time dynamics of the failure rate parameter λ. During degradation, only one parameter of the Markov chain model changes (namely, λ); all other parameters are constant as shown in the figure. For example, the constant nature of parameter μ indicates that CM rates remained the same during the 500-hour observation period. Parameter λ is modelled using the Weibull hazard function with time-dependent parameter β, which plays the role of a degradation variable in this case. The values of parameter β, shown in the bottom panel of Figure 41, are obtained from Equation (3) and represent the degradation index. In Figure 41, the forecasted values of β are plotted for a forecasting horizon of one hour. As seen from Figure 41, straightforward application of the Weibull hazard model produces intuitively paradoxical results. While parameter β began to increase after 200 hours as deterioration of the M&P set progressed, the hazard rate dropped indicating that the M&P set was less likely to fail. Analysis of Figure 38 as well as Equation (12) for the hazard rate reveals the reason for such behavior. For the time period of 0-200 hours, the value of parameter β is slightly higher than 1, which translates into a high hazard rate, as evidenced from Figure 38, in which hazard rates for β=1.1 are significantly higher. When parameter β starts to expand at 200 hours; however, the hazard rate λ decreases as evidenced from the top panel of Figure 41. To avoid this situation, the Weibull model must be modified into a proportional hazard model that can be expressed as follows [17] [18]:

$$\boldsymbol{\lambda(t|X) = \lambda_0(t) \cdot e^{\theta \cdot X}} \qquad \textbf{(17)}$$

where λ_0 (t) is the baseline hazard and X is a degradation variable. The baseline value of λ_0 (t)= $3.6{\cdot}10^{-4}$ was used in this simulation. Parameters of the proportional hazard rate model for parameters λ_0 (t)=$3.6{\cdot}10^{-4}$, μ=$1.8{\cdot}10^{-2}$, η=$6.4{\cdot}10^{-5}$, and ν=$7.5{\cdot}10^{-2}$, are shown in Figure 42. The plots are similar to Figure 37, the top panel shows the time evolution of the parameter λ using the proportional hazard model from Equation (17) while three other Markov model parameters (e.g., μ, η, and ν) remain the same and constant. However, in this case the variable β plays the role of degradation variable (i.e., X in Equation (17)).

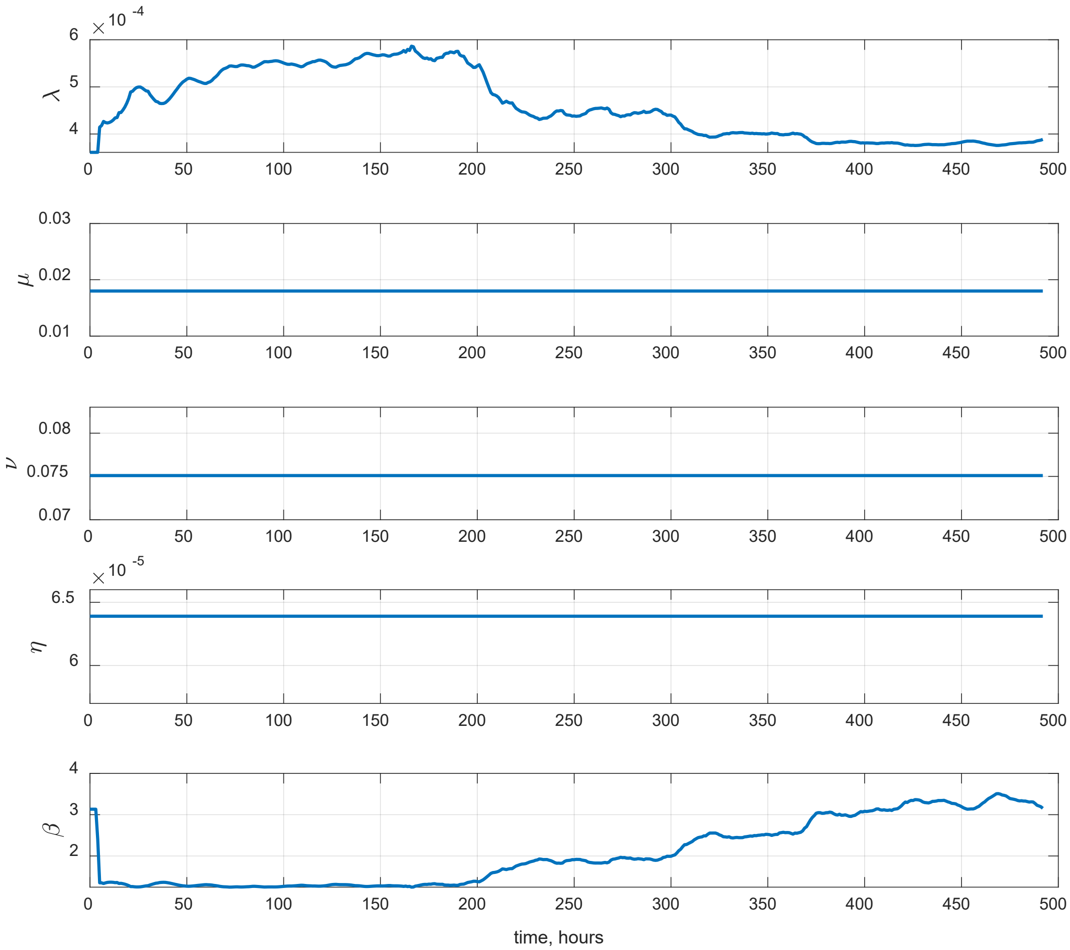

Figure 41 - Parameters of Markov Model for M&P Set Degradation Along with Degradation Variable β using Weibull Hazard Model.

When using the proportional hazard model, parameter λ exhibits intuitively correct behavior as it starts to increase along with increase proportionally with the degradation variable. Figure 43 shows the time evolution of the probabilities for different states of the Markov model when applying the proportional hazard model.

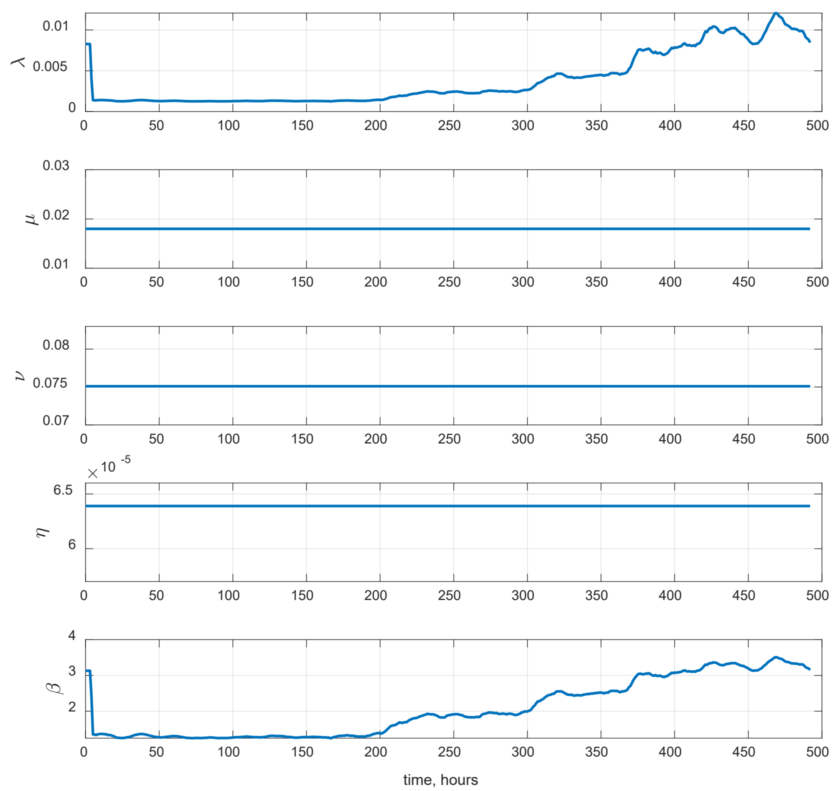

Figure 42 - Parameters of the Markov Model for the Proportional Hazard Model

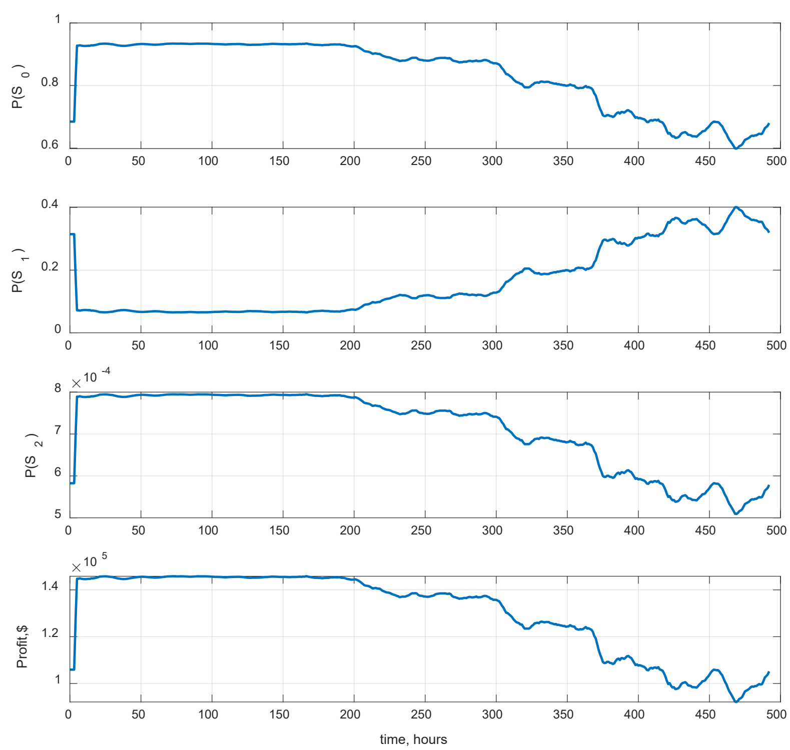

Figure 43 - Markov Model Probabilities of States for the Proportional Hazard Model

As seen from Figure 43, for the proportional hazard model, the probability of being operational $P(S_0)$ decreases as M&P set degradation progresses. On the other hand, the probability of being in CM $P(S_1)$ increases, and that probability of being in PM $P(S_2)$ decreases. This model behavior is intuitively correct, since we expect the M&P set deterioration to eventually lead to repair. Furthermore, the bottom panel shows how the expected profit value decreases as a function of different probabilities of states.

With the probabilities of states for the homogeneous model, and interpreting them as a percentage of relative time spent in a given state, the following cost-benefit model can calculate the expected hourly profit for a 1200 MWe unit with one of its M&P sets down:

$$\boldsymbol{Hourly\ Profit = Hourly\ Revenue\ at\ Full\ Power \cdot P(S_0) -}$$
$$\boldsymbol{\begin{pmatrix} Hourly\ Labor\ Rates + Hourly\ Foregone\ Revenue + \\ Hourly\ Material\ Cost \end{pmatrix} \cdot P(S_1) -}$$
$$\boldsymbol{(Hourly\ Labor\ Rates + Hourly\ Foregone\ Revenue + Hourly\ Material\ Cost) \cdot P(S_2)} \qquad \textbf{(18)}$$

It should be emphasized that while economic analysis of system performance is beneficial in terms of foreseeing economic losses and gains, it can only be meaningfully applied in the case of long-term operations (e.g., for the duration of the fuel cycle). In a quickly developing situation such as the M&P set degradation described above, the profit calculations are only useful as an indicator of how rapidly the system is moving towards economic losses. While the analysis below is similar to the economic analysis presented in [13], it has a significant modification as it includes the total hourly cost of PdM ownership. The total hourly cost of PdM ownership can be significantly higher than simply PdM cost. Also, a significant difference from the previous analysis is consideration of the cost of false alarms, which was not presented before.

The economic analysis can also be applied to calculate the potential benefits of introducing a new on-line PdM system, as it enables estimation of its economic benefits. In this case, the hourly profit equation can be modified as follows:

$$Hourly\ Profit = Hourly\ Revenue\ at\ Full\ Power \cdot P(S\_0\ ) - (Hourly\ Labor\ Rates + Hourly\ Foregone\ Revenue + Hourly\ CM\ Material\ Cost) \cdot P(S\_1\ ) - (Hourly\ Labor\ Rates + Hourly\ Foregone\ Revenue + Hourly\ PM\ Material\ Cost) \cdot P(S\_2\ ) - \boldsymbol{Total\ Hourly\ Cost\ of\ PdM\ Ownership}$$

$$\boldsymbol{Total}\ Hourly\ \boldsymbol{Cost\ of\ PdM\ Ownership} = (\text{Capital expenditure + operational expenditure})$$

Assuming that the price of the PdM=135K, the hourly O&M labor \$100, and the hourly operational cost of materials is \$333, then the hourly cost of PdM for 9,000 hours of operation is 135K/9000=\$15.

For example, for Salem Unit 1 – using parameters $\lambda=3.6\cdot10^{-4}$, $\mu=1.8\cdot10^{-2}$, $\eta=6.4\cdot10^{-5}$, and $\nu=7.5\cdot10^{-2}$ – the probabilities of the three states are p_0 =0.97952, p_1=0.01964, and p_2=0.00083. The PdM system does not directly affect the failure/downtime rate, though it does affect the CM and PM rates, as it enables both to be performed quickly and efficiently. Figure 44 shows the change in profit after simply introducing PdM and assuming that the PdM increased the CM rate – μ by 20%, and PM rate -ν by 10%. The PdM was put into operation at 3,000 hours. As is seen, simply introducing PdM without any change in rates leads to a reduced in hourly profit, since the utility must absorb the cost of PdM ownership. However, even with a 20% increase in CM rate and a 10% increase in the PM rate, PdM will justify the spending and increase the hourly profit by \$79. This example demonstrates how utilities can use the proposed economic modeling to make decisions on introducing new technologies or implementing new maintenance policies. For example, prior to purchasing a new PdM system the utility may want to investigate how much it will affect its maintenance rates and whether the benefits will justify the spending.

Another aspect to consider while contemplating the use of a PdM system is the rate of false alarms generated by a PdM systems, which are notorious for false alarms. If we assume that a PdM system has a relatively low false alarm rate of one per year (φ= 1/8760 hours=$1.14\cdot10^{-4}$ 1/hour), and using the above-described rates for Salem Unit 1, we can estimate how much a single annual false alarm will affect the system's economic performance. False alarms directly affect the PM scheduling rate in the Markov model -η, as each false alarm initiates a trip to the PM state in the Markov model. Assuming $\eta = 6.4 \cdot 10\text{-}5$ we must add the false alarm rate φ to η in order to obtain the false-alarm-adjusted η. The results are shown in Figure 45.

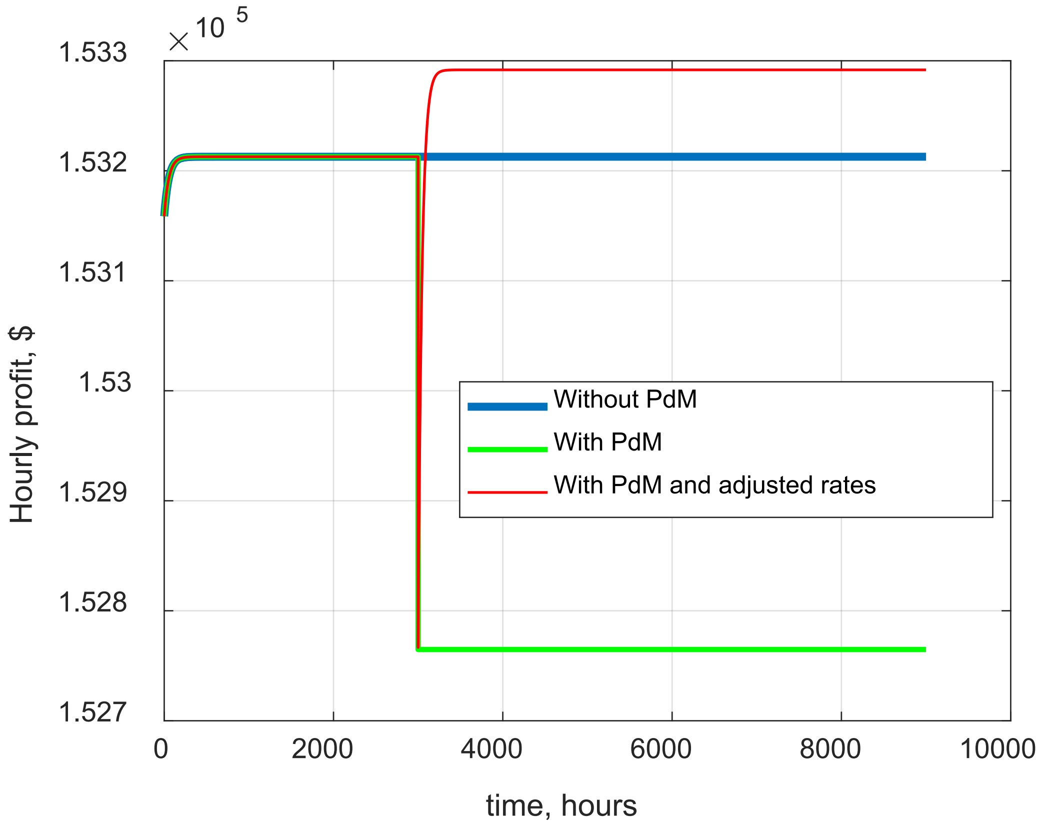

Figure 44 - Change in Hourly Profit After Introducing PdM and with Change in μ by 20% and ν by 10%

As evident from Figure 45, a single false alarm makes introducing PdM detrimental to the unit's economic performance, as hourly profit drops by $14. In general, however, a PdM system should reduce the failure/downtime rate λ, as the system will be directed into a PM state once degradation is detected. If significant enough, this reduction in λ may provide economic benefits, as shown in Figure 46. Figure 46 shows the change in hourly profit for a 10 % change in λ, a 20% change in μ, a 10% change in ν, and a 78% change in η due to false alarms. We can see that, in this case, the hourly profit is increased by $176, making PdM a worthwhile investment.

In summary, we presented a Markov chain model with a proportional hazard model and demonstrated how it can be applied to forecast the probability of different states during M&P degradation. We also demonstrated how an economic model can be used to calculate economic benefits and losses for different maintenance scenarios.

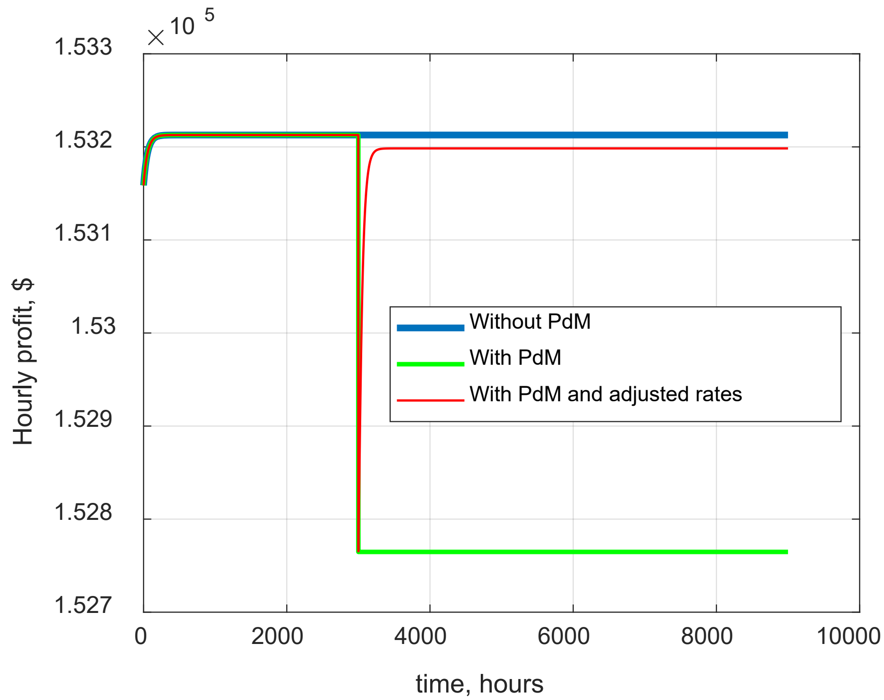

Figure 45 - Hourly Profits for the System Without PdM, with PdM and No Rates Adjustment, and PdM with False Alarm Adjusted Rates

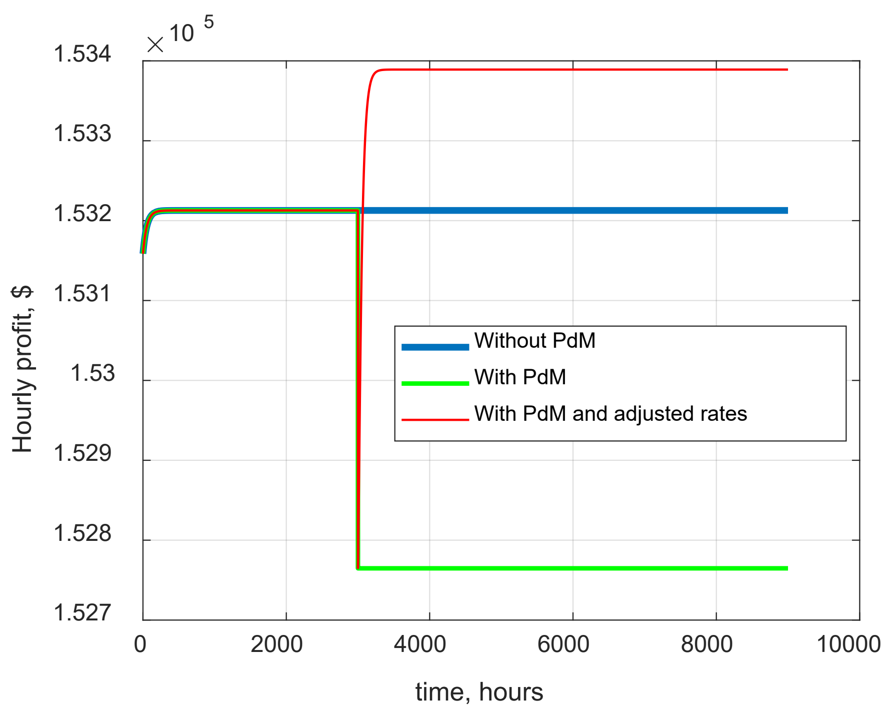

Figure 46 - Hourly Profits for the System Without PdM, with PdM and No Rates Adjustment, and With False Alarms and λ Adjusted Rates

5. ADVANCED PATTERN RECOGNITION MODELING

Another approach to CBM is to use Advanced Pattern Recognition (APR) data – enhanced with engineering features – to identify the state of components in the field. PKMJ worked with APR software to provide a comparison for other CBM work performed on the Salem CWS components. This approach allowed the project team to see how a mature modeling software analyzed data associated with the Salem CWS, as well as another option for alarms to be integrated into the NDP.

5.1 Fault Analysis Scope

The first step in performing analysis using the APR method is to evaluate the scope of review for the analysis. In this case, the focus of the project is on the major components associated with the Salem CWS in support of PMO work performed during Goal 1. The APR modeling scope included two pumps and motors each, as well as the downstream condenser, as interactions between the pumps and motors were thought to add value to the modeling. For example, pumps 11A and 11B share a common discharge header on the outlet side of the condenser. This impacts the water outlet temperature monitoring for discharge back into the river. Other shared or associated parameters were considered and included to provide a holistic view of system and component performance within the model.

Modeling efforts for plant components must consider whether the scope of review is limited to component-specific events as opposed to both component and system events. This decision governs the total amount of system parameters that needing consideration within the model. Experience from this effort shows that including system parameters within CBM models provides system-level details on trends that would not be identified at the component level. An example of this is when a pump startup on one train of the Circulating Water System resulted in a temporary vibration alarm on another pump due to their relative location within the building allowing some interaction. Understanding the direct and indirect interactions among components within a monitoring scope is a good practice for generating better insights.

APR Model development also required the selection of an optimal dataset as review of the entire component history from the beginning of the available plant process data (starting from 2010) would present data processing and evaluation challenges. As a result, PKMJ opted to focus on a smaller dataset to ensure that model performance is maintained. In addition, the baseline for the component is reset after each major replacement or maintenance task. Using a large dataset that bridges across those major replacement tasks may provide less-than-ideal results for the APR modeling.

Since a review of the data for pump and motor sets 11A and 11B showed a significant change in the system's baseline performance following the Spring 2019 outage, that outage was selected as a starting point. The end point of the analysis was January 2021, aligning with preliminary analysis results from the other CBM modeling method discussed via this report and freezing the analysis scope. This timeframe selection was also supported by the availability of vibration sensors installed after September 2019, as previously presented in a Technical Report for this project [2]. The availability of vibration data unlocked additional engineering features that could be reviewed in combination with the available plant process (PI computer) data features.

For long-term modeling performed using live data, establishing the baseline is less critical, as streaming data is configured to continuously re-calibrate the model. In this context, "live data" refers to real-time or near-real-time data obtained, at maximum, every 15 minutes. It is highly recommended that live data be utilized in modeling solutions, as live data allows for modeling real-time conditions and proactive response to degradation risks. The approach of using historical data was applied because it was unknown whether a live failure would occur during the period when this model was being developed. As a result, to demonstrate the viability of the technology, the team utilized and tested their models against historical data that featured known results. This also allowed for a comparison to the results generated by INL as described in Section 4 and the previous Technical Report [13].

The next scoping decision was to determine which faults of interest provided the most direct value in terms of monitoring component health. PKMJ utilized feedback from PSEG to identify the faults of interest. The component and system faults identified for review are as follows:

- Diffuser Failure
- Waterbox Fouling
- Dirty Motor Intake Screens
- Motor Failures (including Moisture & Salt Contamination and AC Motor Faults)
- Bellmouth Failure
- Misalignment
- Bearing Fault
- Pump Failures (identified primarily by Vibration)

Real-life examples were used to the greatest extent possible when evaluating these faults which could lead to potential failures. However, the faults of interest included event types that did not occur between 2019 and 2021. The remaining faults were configured for simulation within the software. This approach was selected to show that simulated data could be used to identify events through the analysis, despite the fact that the events did not occur during the monitoring period.

5.2 Data Alignment to Fault Characteristics

Identification and interpretation of component and system faults through APR requires engineering analysis of the associated data and system design. The team worked with subject matter experts (SMEs) to configure models based upon the plant configuration. The engineers working on the project received input from personnel who had worked with vibration data for various component types. This knowledge allowed the team to establish a profile of anticipated responses for each fault type. These fault profiles were then modeled in the APR software as the signal responses triggered by signal- or feature-specific fault detectors. A summary of example data types associated with the primary faults of interest are listed in Table 9.

Diffuser Failure
- Vibration (at multiples of 6x motor running speed or 29.4 Hz)
- Motor Current Decrease
- Stator Temperature High
Waterbox Fouling
- Motor Current Increase or Decrease
- Inlet Pressure Increase
- Waterbox Delta T (Change in Temperature)
- Gross Load Decrease
- Vibration (approx. 30-40 Hz per experimental data)
Dirty Motor Inlet Screen
- Motor Bearing/Stator Temperature Increase
Misalignment/Unbalance
- Vibration (at 1x or 2x running speed or 4.9Hz and 9.8Hz, respectively)
- Motor Bearing Vibration
Motor Air Intake Clogging
- Stator Temperature Increase

Bearing Fault
- Vibration (at motor running speed or 4.9Hz)
- Bearing Temperature Increase
Motor Failure
- Motor Stator Temperature Increase
- Vibration (at 2x Line Frequency or 120Hz)
- Bearing Temperature Increase
- Vibration (Magnitude of peak acceleration)
- Inlet Pressure (Hi)
Bellmouth Failure
- Motor Current Decrease
- Waterbox Delta T (Change in Temperature) Increase
- Pump Vibration
Pump Failure
- Vibration (Magnitude of peak acceleration)

Table 9 - Failure Mode Characteristics

These different characteristics provide a wide scope of parameters usable for identifying events that impact component health of the CWS pump and motor, as well as system health, as the pumps and motors are impacted by events associated with the condenser waterbox. By reviewing these primary parameters associated with each fault, the team determined which other secondary parameters were correlated with each primary parameter. Primary parameters are the main characteristics of a fault scenario and secondary parameters are considered plant process variables that correlate with the primary parameters. Evaluation of the secondary parameters provided guidance for substituting other available data when primary data sets were unavailable and provided insights into other characteristics used to identify the fault scenarios of interest. The data sets were analyzed by the APR software to produce a correlation report showing which parameters were most closely related. This report was used to group the parameters within the project, which is a key element of the APR analysis used in identifying the appropriate models.

5.3 Data Interface with the Model

Another critical aspect of performing advanced analytics is ensuring that the data are available for the modeling tool to utilize. This includes placing data in an accessible location and formatting them to be utilized by the software. This process introduces some challenges as the datasets are generated by various sources (PI computer, vibration sensors) and lack consistent formats. PKMJ pre-processed the raw process data to convert the datasets from the source formats into a consistent format for use by the APR software.

Data pre-processing is important when developing advanced analytics and integrating it into a digital platform. The pre-processing must consider the required inputs and outputs of the model to support data utilization and the prevention of any data modification that could impact the results. In this context, data modification, includes simplifying the data by altering them in such a way that modifies the original results (e.g., averaging to make sampling frequencies match).

Advanced modeling should utilize data sampled at ten times the desired response of the system. This allows the model to analyze changes in condition and identify alarms in a timely manner. A single data point does not generally provide sufficient evidence of an alarm, but a change registering across five to ten datapoints can establish a trend that is acknowledgeable by the model. For example, raw vibration data sets can be sampled several hundred times per second; however, the plant process data used for this project was retrieved once every hour. This led to several modeling challenges, as the datasets were inconsistent with each other. Addressing this issue necessitated, addressing the model's responsiveness in advance so as to understand the utilization of outputs in downstream processes.

For the scope of this effort, identification of emergent errors in a period less 12 hours was deemed a low priority. The project goals focused on generating degradation risk insights proactively, with sufficient time to address the risks in a cost-efficient manner. Data sets obtained on a per second or per minute basis are more frequent than required to meet this project goal. The team decided to baseline all data on an hourly basis, with data obtained more frequently being averaged. This approach must be adjusted for data sources centered around zero (e.g., vibration, which can be positive or negative). This issue was resolved by taking the absolute maximum value of the vibration data every five minutes and then obtaining an hourly average of those maximum values.

A future consideration for any type of CBM or APR modeling within digital platforms and services is to understand the required responsiveness of the model to plant events. If a plant wishes to monitor components to identify emergent issues and respond within a day or less, the associated data structures, processing, and tools must be defined to suit that approach. This includes the amount of data required to train a model, as modeling can be performed to show short-term or long-term trends. Utilities and vendors developing a digital platform should consider these downstream service requirements when designing their data architecture and interfaces.

Parallel to the discussion of data sampling or data frequency, it is important to understand whether the model must respond to data in real-time. Two potential approaches for modeling include performing bulk data analysis once per selected period (i.e., a day) or performing real-time analysis to immediately identify issues. This decision impacts the results of analysis models and how they benefit the plant. Bulk data analyses performed on a set frequency may support long-term maintenance and monitoring goals but not address emergent work identification. However, bulk analyses may reduce the computation costs required to provide utilities with solutions of value. Inclusion of this consideration in design requirements is another critical aspect of supporting the underlying platform and its services.

5.4 Modeling Approach

After verifying the scope of the model, confirming which faults were of interest for the CWS pump and motors as well as the pre-processing of the required data, the first major task was to identify the system's baseline performance. Calibrating the APR software to understand what normal condition data looks like for the target system (i.e., healthy state of the CWS pumps and motors). The calibration dataset is established by removing bad data, such as data obtained while the system is offline or during sensors malfunctions (i.e., failed high/low or offline). Another similar issue was that the vibration data sampling frequency was inadvertently adjusted by the vendor during the project timeline, resulting in problematic data outputs until the sampling frequency was restored.

If the removal of bad data resulted in insufficient data to establish a reliable baseline for a given parameter, the parameter was removed from the model to allow the remainder of the model to perform as intended. This occurred for the 11B diffuser event, as the vibration sensors were non-functional for the pump during the event. It is recommended that the data removed from the model be minimized to prevent introducing undue bias into the model results. Where possible, secondary indication parameters were included in the model when the primary parameters were insufficient to establish a baseline. For example, as the 11B pump vibration sensors were out of service for an extended period due to battery issues, the 11B motor vibration parameters were substituted into the model as the vibration translates to a degree through the interface between the pump and motor. This is confirmed via correlation of these parameters, as was discussed in Section 5.2.

Use of secondary parameters was implemented as a compromise between using the primary (or ideal) parameters and still considering the real-life issues with maintaining instrumentation, including temporary sensor outages or failures. The model becomes more diverse and robust as complimentary sensors included in the model, which improves the predictive value of the model. Installation of additional sensors must be balanced with the cost of installing and maintaining those sensors. However, when redundant sensors are available, they should be considered for inclusion in CBM models. It is also critical to understand the consistency or correlation among these parameters to ensure that the models are not poisoned by the inclusion of additional features. It is highly recommended that a correlation threshold be established for utilizing secondary parameters to ensure that the model maintains its accuracy and sensitivity to events. Without implementation of a threshold, poorly correlating sensors may be substituted within the model, which will result in less-than-ideal model outputs.

Also critical for evaluating the model data is understanding the plant's overall condition. The CWS is the direct interface with the Salem NPP's ultimate heat sink (i.e., the Delaware River). When the plant is not operating, the CWS is not demanded to perform to the same requirements as compared to full-power operations. Based on this understanding, to evaluate the performance required of the CWS, the model utilizes gross load for the plant as a factor in determining when the plant is running at, or near, full power. This same consideration is included for the circuit breaker position for the CWS motors, as the system responds differently when the individual pump and motor sets are removed from service. APR model design and development should consider configuration requirements at the plant, system, and component-level to obtain the optimal results and insights.

Once the model is calibrated, testing datasets must be created to model system events (or faults). Test datasets were created from the historical data by segregating known component issues such as the 11B diffuser event from April 2020 and several instances of waterbox fouling events. These events were placed into test datasets to evaluate how the model responds to frequent or sudden system events. When selecting the test datasets, it was critical that a majority of the sensors were active during the period of interest, and that potential faults could be matched to historical WO data to provide a reference in support of the evaluation of the model's results.

APR models take advantage of the calibrated data to establish a predicted baseline that can then be compared to the observed values to determine whether parameters are abnormal. Each signal being evaluated by the model is placed in a fault detector of the APR software to identify whether the difference between the predicted and observed (residual signal) values exceed a defined threshold. This method allows the model to anticipate conditions that impact parameters during normal system operation: high tide levels, ambient temperature, etc. As noted above, utilities typically generate alarm conditions based on parameter setpoints, which are fixed values. Use of the residual signal is more robust for complex designs for which multiple factors contribute to the parameters. Setpoints established by utilities are highly conservative compared to system performance. In addition, these APR fault detectors are configured to trigger on a delay, thus minimizing nuisance alarms by ensuring that the condition is confirmed prior to generating an alarm.

Once the APR fault detectors are developed, they can be associated with the fault parameters discussed in Section 5.2 to provide the basis for a diagnosis. The grouping of these various parameters must be well understood prior to creating this type of model to assign physical meaning to the data being analyzed. At a base level, APR models do not understand engineering concepts, so the model produces alarms based on whether signals deviate from their predicted values. SME's or modelers with technical expertise must configure the pattern recognition outputs to assign a diagnosis to a selected set of alarm conditions to produce a desired diagnosis output. A major consideration is how the model can determine the physical problem so that data inputs and data processing can be aligned to support the model in making correct diagnoses or decisions.

The primary signals associated with the diagnosis are then entered into a prognosis model to determine the time until components reach an alarm threshold, thus requiring the component degradation to be addressed. Using the prognosis model outputs provides guidance on how long components will last before CM is required. A prognosis output is critical for scheduling and planning work required to address an alarm, as the prognosis output of a model can be used to estimate when craft teams must be available and when associated stock must be on-site to support maintenance.

5.5 Modeling Results

The scope of the modeling focused on faults of interest associated with the Salem CWS. The primary degradation events were compared to the results presented in Section 4 and the previous Technical Report [13]. The faults examined in this section are the 11B pump diffuser as well as the 11A and 11B condenser waterbox fouling events. These issues were focused on because they provided a sampling of issues that occurred during the project while the vibration sensors were active. This section will review the model outputs for each event and provide a point of comparison between the two different model approaches.

5.5.1 CWS Pump 11B Diffuser Event – April 2020

Salem's CWS had one relatively major event occur during the period in which the vibration sensors were installed. In April 2020, the 11B CWS pump experienced a failure of the diffuser associated with the pump. Post-failure reports from walkdowns and inspections indicated that the diffuser had separated from the propeller shroud and was deposited at the bottom of the bay (underwater). The mounting hardware holding the diffuser to the propeller shroud was missing, and several bolts were sheared. This event can be considered an uncommon occurrence for the CWS pumps, and was of note throughout this research project because the associated process parameters progressively worsened over the two-week period until the equipment was shut down. This section details the APR software results for this event.

The 11B CWS pump and motor was removed from service on April 21, 2020 after the component was determined to be deteriorating. Site Operations noted a potential issue on the morning of April 20, 2020, after completing maintenance on the 11A CWS pump, with the system responding in a different manner (a higher drop in gross load) during an on-condition waterbox cleaning. The plant generated a notification to troubleshoot a potential issue with the 11B CWS pump and motor at that time, as the issue was still unverified. The team in the monitoring center worked with Operations to determine that several parameters (e.g., current and vibration) were indeed trending towards an abnormal condition. Following that review, the CWS pump was removed from service, and the above-mentioned damage was later identified.

APR software would have streamlined the identification of this same issue and provided additional insights. When focusing solely on plant process parameters from the PI computer, a major indicator of the diffuser issue was the decrease in the 11B motor current. The APR software generated an alarm as the observed value of the current showed a sustained trend below the expected value, starting on April 18, 2020. Figure 47 shows that the residual (the difference between the observed and expected values) for the 11B motor current began an identifiable downwards trend on approximately April 17th, but this trend was not significant enough to trip the APR fault detector until April 18th. Based on this, it is reasonable to say that, if utilizing the APR model, the site would have been notified two days earlier of a potential issue associated with the motor, as compared to using the current process for evaluating issues and the data available prior to the beginning of this project.

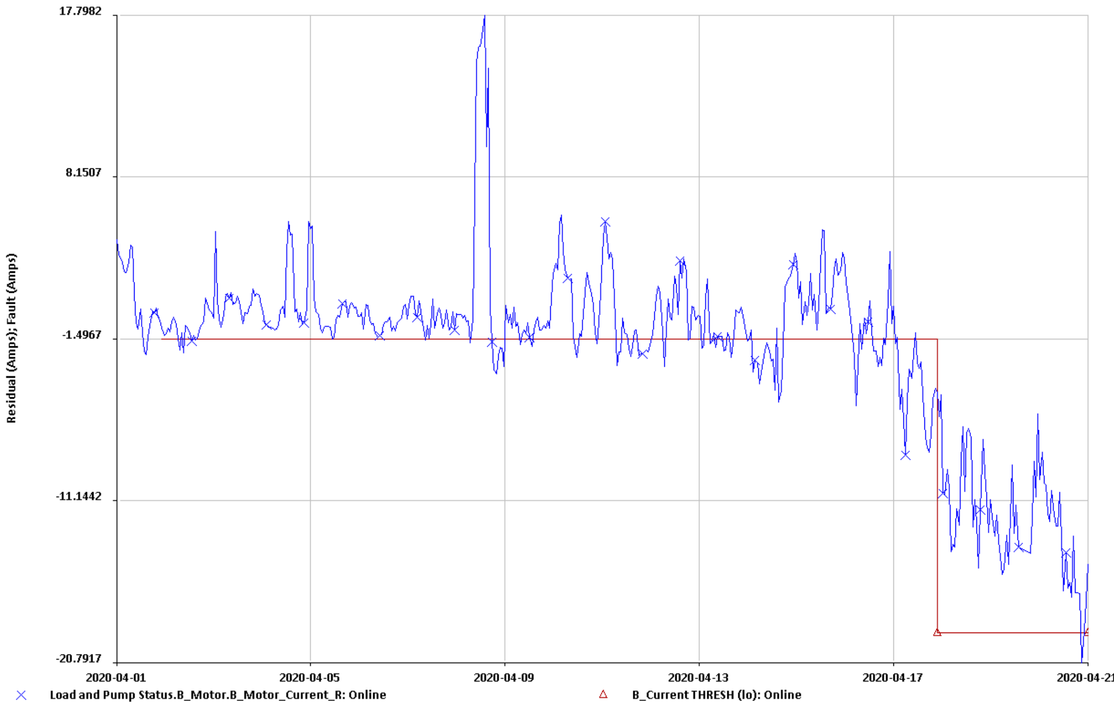

Figure 47 - Residual Graph of Motor Current for 11B Diffuser Event

Phase 1 of this project involved installing vibration sensors on the CWS pumps and motors, as discussed at length in this report. The vibration data sets were also reviewed in the APR software to determine whether vibrations provided any additional indication of the diffuser event. The vibration was configured to alarm within the APR software based how many standard deviations the observed value was removed from the calibrated (expected) baseline. Ideally, the pump vibration would be used in this case; however, the pump sensors were out of service during this event. Based on the correlation analysis described in Section 5.2, the APR software was configured to use motor vibration as a secondary parameter when the pump vibration was unavailable, as the two are highly correlated. As seen in Figure 48, motor vibration indicated an alarm state on April 10, 2020. This alarm was generated a full 10 days prior to the above-described troubleshooting actions taken by site Operations.

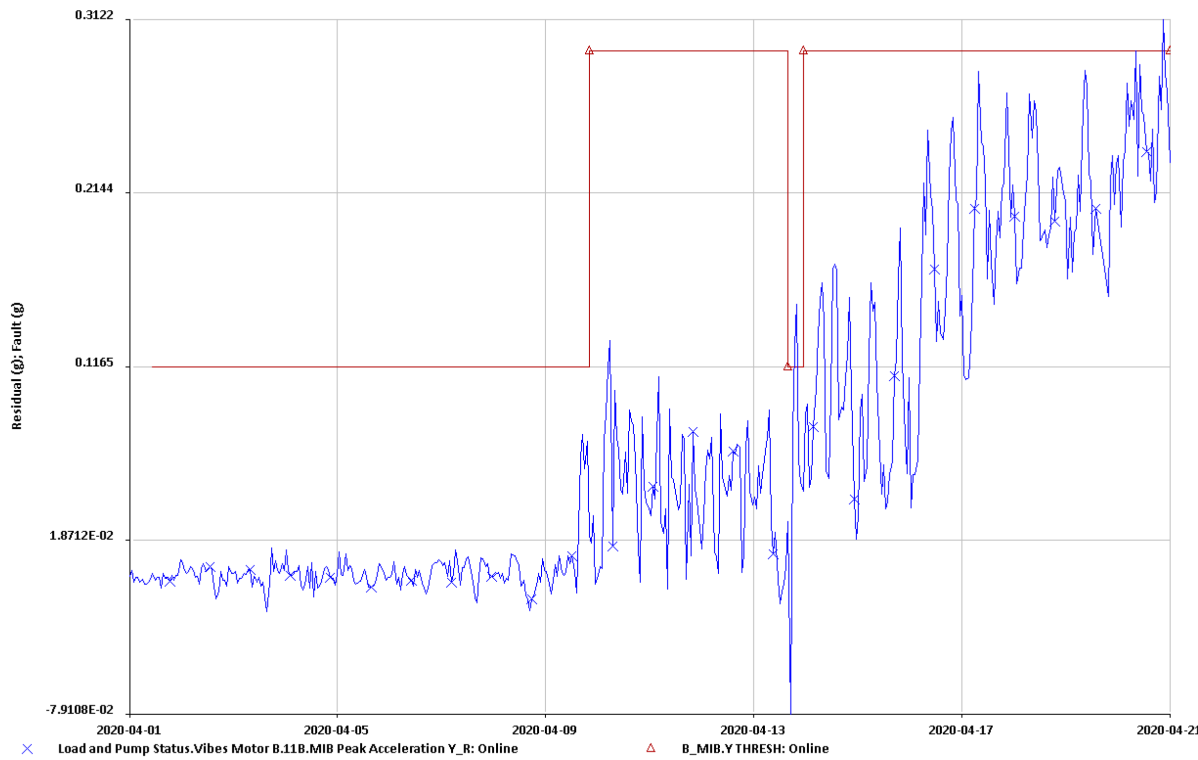

Figure 48 - Motor Vibration (11B MIB) Associated with the 11B Diffuser Event

During the alarm condition, the vibration consistently trended upwards, which would have provided Operations and monitoring personnel with additional confidence that the 11B pump and motor set had a condition that was deteriorating – well before the plant process data indicated an issue. These results occurring in real-time would have provided PSEG with immediate insight into an issue requiring additional investigation. This could have been combined with the NDP automated work management interface (see Section 6) to prepare a WO for removing the 11B pump and motor from service in order to perform repairs.

5.5.2 CWS Pump 11B Waterbox Fouling

Salem's CWS has frequently experienced condenser waterbox fouling or grassing events caused by debris from the Delaware River entering the CWS. When debris enters the CWS, some debris is deposited in the system, inhibiting system performance. Specifically, Salem is concerned about a decrease in the plant's thermal performance, as excessive fouling limits flow through the condenser. Decreased flow through the condenser changes how much heat is rejected from the plant – a critical characteristic for maintaining optimal plant performance.

The waterbox fouling event discussed in this section occurred in December 2020. The best indication of waterbox fouling is the difference between the CWS inlet and outlet water temperature. The inlet water temperature is gauged in the Circulating Water Building bays and the CWS outlet water temperature is gauged at the pipes sending water back into the river. The residual for the change in temperature (or Delta T) began to trend slightly upwards on December 25, 2020 as shown in Figure 49. The first APR software alarm for Delta T was identified on December 29th at 22:00, just before the pump and motor set was removed from service. Also shown in Figure 50 is the motor current residual, which is another common indicator of issues related to waterbox fouling. PSEG uses impact to thermal performance as their key decision point for removing the pump and motor from service to perform on-condition waterbox cleaning. Thus, this result is reasonable.

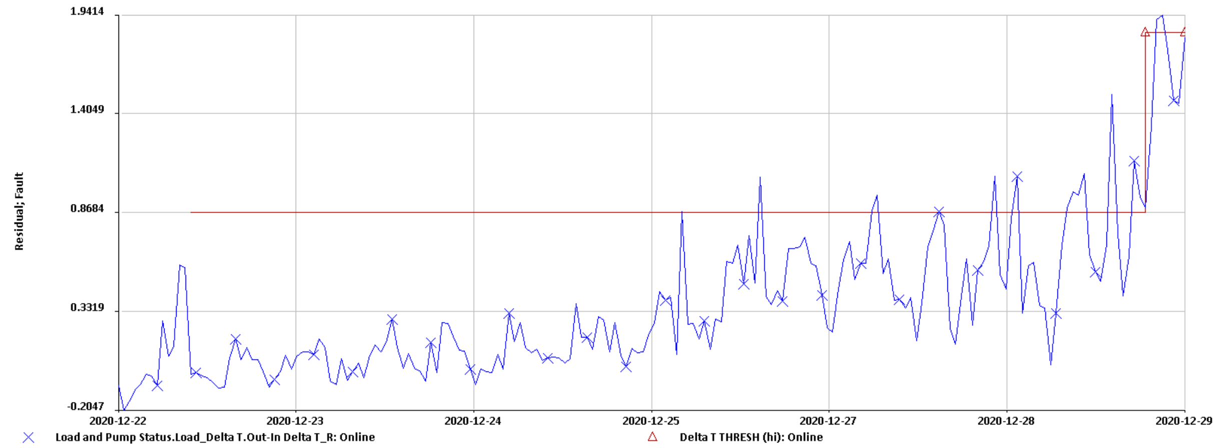

Figure 49 - 11B Waterbox Fouling Change in Temperature (Delta T) Residual

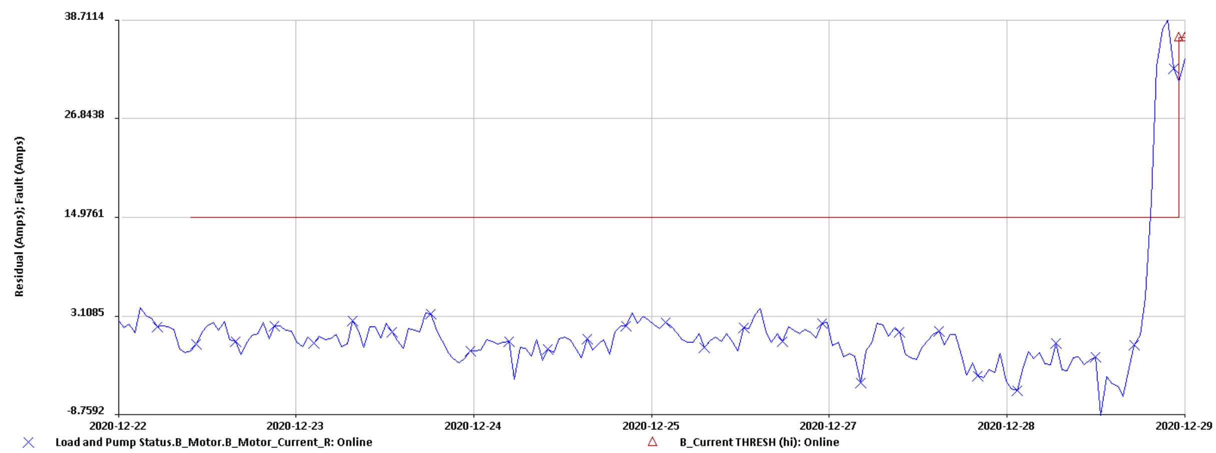

Figure 50 - 11B Motor Current Residual During Waterbox Fouling Event

When incorporating vibration into the APR software for this event, an alarm condition for the motor was identified on December 24, 2020 (see Figure 51). The motor vibration alarm was developed using a combination of each motor vibration sensor and their outputs in the X- and Y- directions. The alarm for the motor indicates that multiple vibration parameters exceeded the alarm threshold (again, by evaluating how many standard deviations the observed data was removed from the expected observation). As a result, the site would have been able to compare the excessive vibration with the trend in Delta T to obtain a prediction an additional 3 days before thermal performance was impacted by this waterbox fouling event.

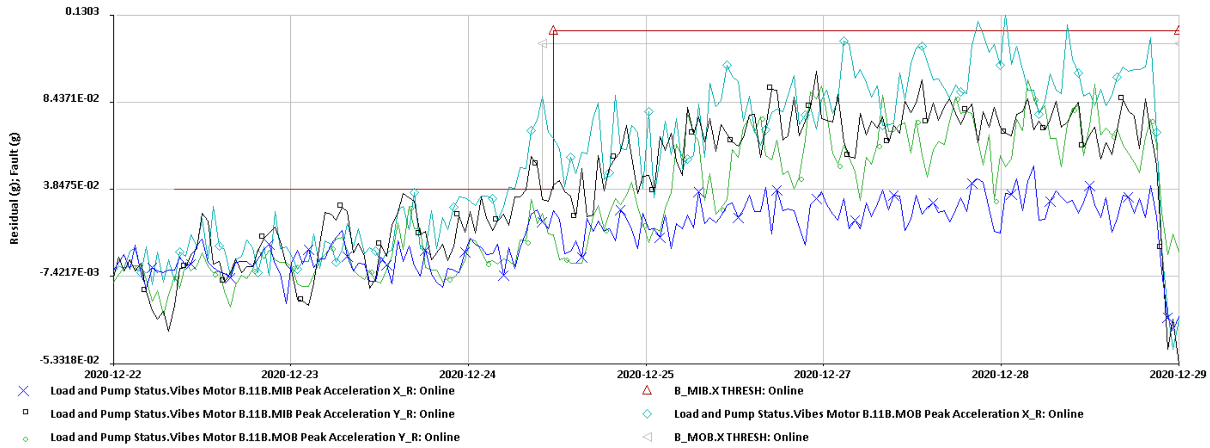

Figure 51 - 11B Motor Acceleration for Waterbox Fouling Event

Another insight of interest for the waterbox fouling event, is that the waterbox fouling for Salem's CWS tends to show an increase in the frequency response between 30 - 40 Hz. Although this was determined experimentally through various examples reviewed during this project, this insight into the data provided another process characteristic that could be included within a degradation risk model. When using vibration data or other data obtained with a high sampling rate, it is highly recommended to evaluate the data's frequency response across all frequencies to determine whether new features can be developed to better identify specific fault signatures.

5.5.3 CW Pump 11A Fouling

As waterbox fouling was the most frequent occurrence during this project, another event classified as a waterbox fouling event was reviewed for comparison to the one reviewed in the previous section. This event occurred for the 11A pump and motor set on April 19th, 2020. This event is of particular interest, since it occurred during the same time period as the diffuser event described in Section 5.5.1. As noted previously, for pumps 11A and 11B, Salem's CWS utilizes a common discharge header for water returning to the river, so this example was considered to represent the possible impacts of system interaction. In addition, this event proves the need to understand and consider alternate parameters for fault signatures as plant impact may vary for a fault type.

The process data parameters from the PI computer (excluding vibration) differed for this event, compared to the profile observed for 11B. Specifically, the Delta T parameter within the APR software alarmed briefly on April 15th but was quickly restored to near-normal conditions until April 19th, when the pump and motor set was removed from service. The overall trend of Delta T is seen in Figure 52. The 11A motor current was the first PI parameter trigger an APR alarm, which occurred on April 14th. This parameter remained in an alarm state until the pump was removed from service on April 19th. The change in motor current over time is shown in Figure 53. In this case, the data could be interpreted to mean that the alternate PI parameters indicating a fouling event away from the condenser drove this alarm. As a result, the fouling may have been cleared from the condenser by normal flow through the system, but impacts can still be seen on the pump and motor set due to the overall change in performance. The process parameters alone do not tell the full story of what occurred in this event.

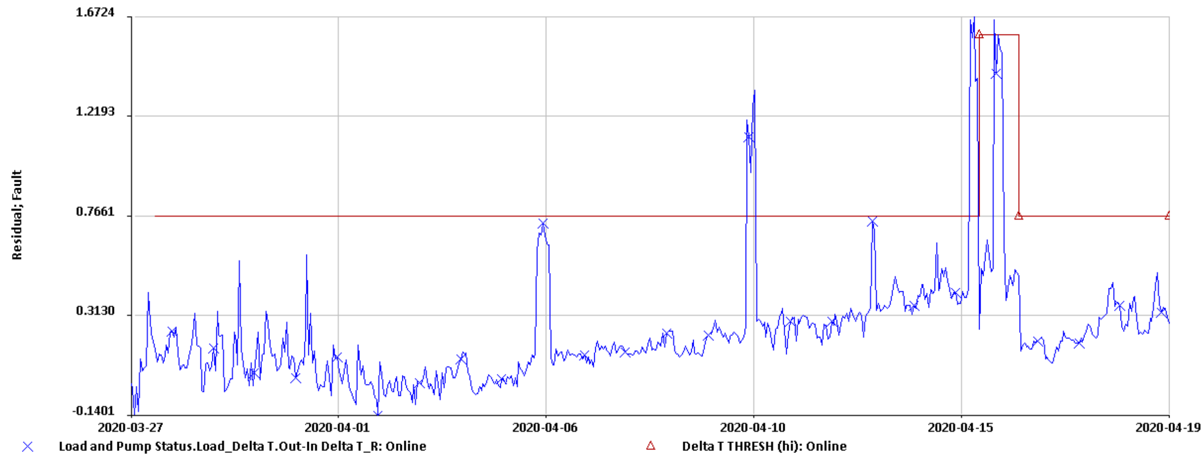

Figure 52 - Delta T for 11A Waterbox Fouling Event

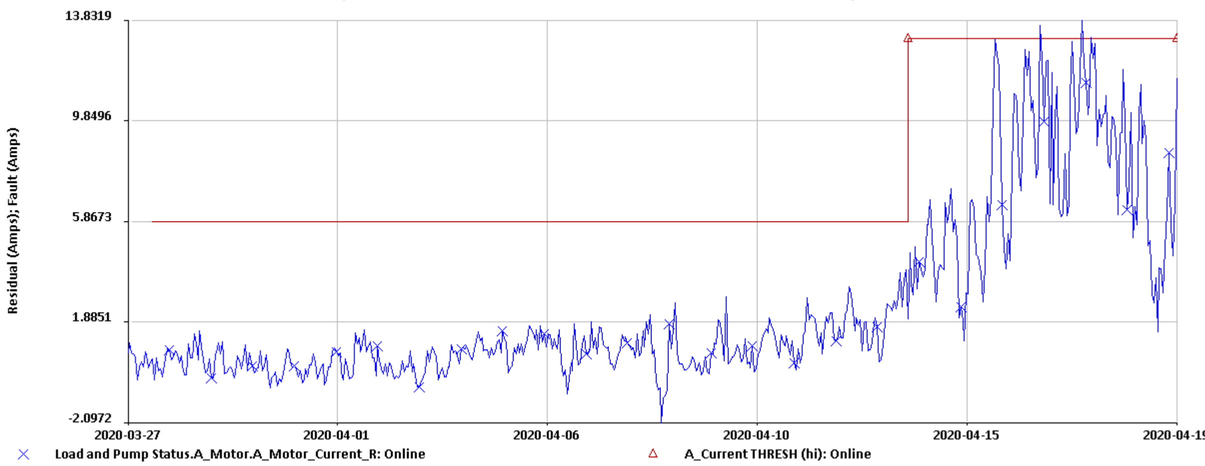

Figure 53 - 11A Motor Current for Waterbox Fouling Event

Based on the above, the same analysis was performed while including the vibration data, which showed an increase in peak acceleration for the 11A MOB vibration sensor at approximately the same time that Delta T was in an alarm state, and that the MIB vibration sensor was in alarm for a significant portion of the time when the motor current was in alarm. The residual values for both sensors were on the scale of 10^{-2} g's. Based on this, it is seen that a vibration change, which could be misinterpreted as a rounding error by monitoring personnel, was alarmed in the APR software. This event was validated by performance of the on-condition WO for removing fouling from the system. No detail from the WO were available to confirm that the fouling was specifically linked to the pump, but the APR software identified another fouling event whose profile differed from to the one previously examined.

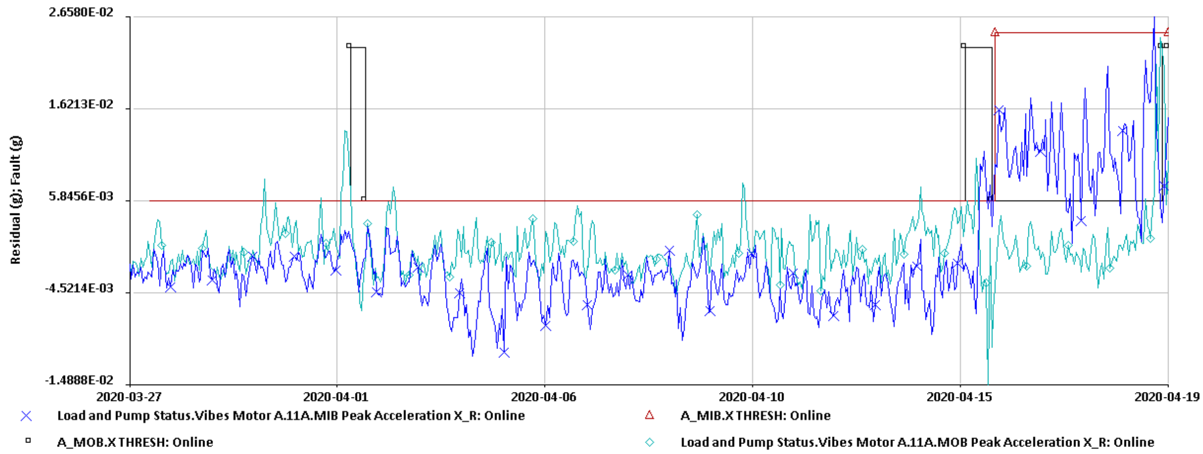

Figure 54 - 11A MIB and MOB Peak Acceleration for Fouling Event

This event for 11A showed that fouling may not be limited to waterbox impacts, thus the fault profile must consider additional parameters when looking for system-level impacts within the model. For simplicity within this research project, PKMJ opted to maintain fouling as one failure mode, but additional review could link these two events to separate diagnoses. With different diagnoses, the response could become focused on where the fouling exists in the system, thus improving the focus and potentially, efficiency of the associated maintenance. As both fouling events had the same overall response (removing the pump and motor from service to eliminate the fouling), determining the exact location of the fouling was not critical. Yet, this example shows the APR software's potential for evaluating plant events.

5.6 Summary

Condition-based monitoring using APR software is a powerful tool for monitoring the condition of components as part of an overall Equipment or Component Reliability Program. APR software was shown to provide detailed insights into events regarding the Salem CWS throughout the period of interest. This tool was enhanced by including vibration sensors installed over the scope of this project to provide enhanced component monitoring capabilities.

A best practice for developing CBM programs is planning for the expansion of the solution to other component types and systems. Each system and component type is unique and requires engineering knowledge to understand how best to model the components and interpret the results. When possible, the engineering knowledge should be "baked" into the results, such that interpretation is minimal or highly structured to enhance the value obtained from the solution. When an alert occurs, the expectation is that personnel can utilize those results as efficiently as possible to support the utility and make decisions regarding maintenance and/or component reliability.

The APR software identified the same events as the CBM model. In addition, it quickly determined when events started and ended – a key element in responding to an event. The APR software required additional configuration so that the model aligned with the pump and motor components, in addition to considering the CWS design requirements. Nevertheless, work performed under this project is easily transferrable to other 4kV vertical pumps and motors with adjustments for system and configuration-specific parameters. The method of evaluating the CWS pumps and motors is scalable to large portions of the plant, including most balance of plant (BOP) and Safety-Related components and systems, with the appropriate rigor applied, depending on the intended application.

6. AUTOMATED WORK MANAGEMENT PROCESS

This section describes the Automated Work Package Process developed in the PKMJ Digital Platform, highlighting the requirements and process-related considerations during development. The goal of the Automated

Work Package tool is to automate as much of the WO generation process as possible to reduce the administrative costs of scheduling and planning work, while still maintaining a high level of technical accuracy. The following key stages were the primary focus areas:

- Utilizing CBM models work to identify component failure modes
- Automating plant stakeholder reviews using data-driven logic
- Linking failure modes to maintenance plans
- Scheduling of work in appropriate work windows
- Generation of WO Packages (Including work instructions, materials, etc.)

PKMJ performed a review of the above focus areas, based on data and procedures available from PSEG, as well as industry guidance from AP-928 [19], in order to identify the most cost-effective value that could be provided without significant modifications to site procedures. An Automated Work Management Process must be accurate, efficient, and minimize administrative burden to be considered for long-term integration to existing processes.

6.1 Work Management Process Considerations

When developing a strategy for work management in the nuclear industry, it is critical to understand the core requirements for an effective work management process. Development of any nuclear work management process would be incomplete without addressing the following key goals and objectives as identified in AP-928:

- Promoting Nuclear Safety
- Improving Industrial Safety Performance
- Improving Radiological Safety Performance
- Improving Equipment Performance and System Health
- Optimization of Safety System and Refueling Outage Durations
- Supporting Effective Station Backlog Management
- Increasing Productivity through Efficient Use of Resources
- Improving Schedule Creditability and Stability
- Reduction of Costs

Based on the above, PKMJ researched existing plant and industry work management processes to define how work management is performed today. Work management includes several phases in which distinct actions are taken, from identifying an issue to, pre-work screening and planning to, work execution and post-work analysis. The high-level listing of each phase and the sequence in which they are performed is shown in Figure 55.

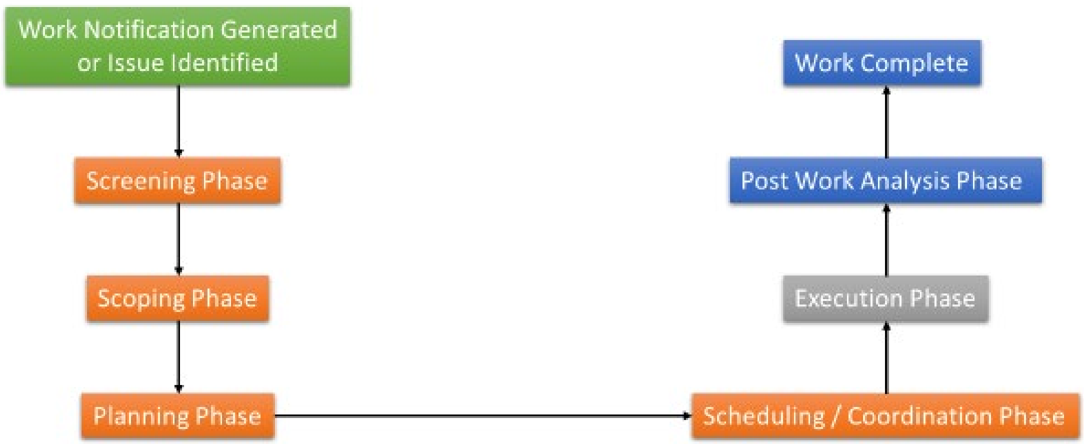

Figure 55 - Legacy Work Management Process

This report will focus on the above noted key categories and address procedural and process steps that represent the most cost-efficient value for implementation of an Automated Work Management Process. Figure 56 shows the revised process flow for work issues entering the automated system. The intent of the developed Automated Work Management Process is to mirror the utility's existing work process, while also utilizing AI/ML techniques to improve efficiency and accurately resolve plant anomalies.

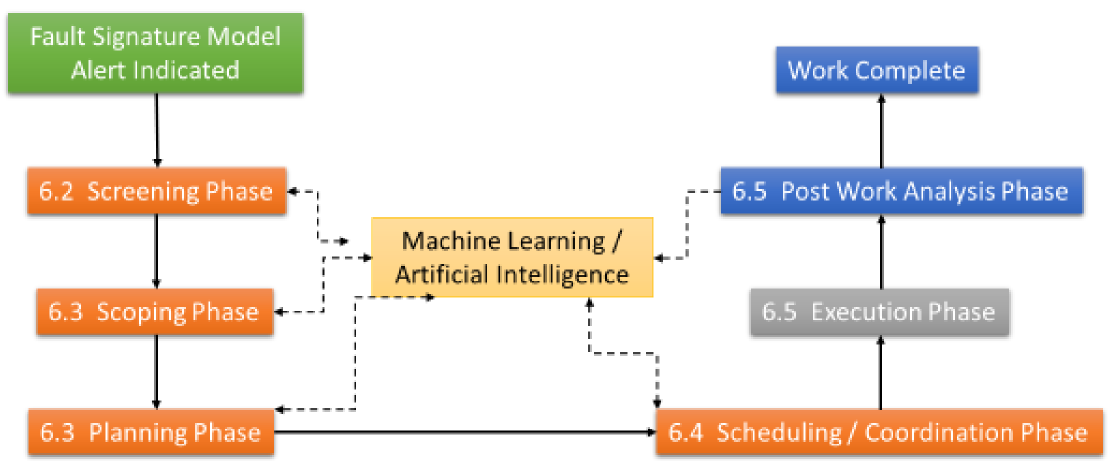

Figure 56 - Automated Work Management Process

6.2 Work Screening Phase

The work screening phase of a Work Management Process identifies the type of work required and gives initial estimates related to work complexity and the materials required. This initial screening determines work complexities ranging from the simplest (i.e. using a wrench to tighten a bolt) all the way through to the most complex (i.e. performing a major plant modification).

The work screening phase focuses on data-driven decisions happening automatically based upon output of the Fault Signature Model and underlying utility enterprise data. Integration with CBM solutions allows for the generation of failure risk insights that can be associated with corrective maintenance for the selected components. The ability to proactively generate a WO based on the risk of a failure as opposed to simply reacting to a failure minimizes the duration and extent of maintenance. The process flow for the Screening Phase is shown in Figure 57.

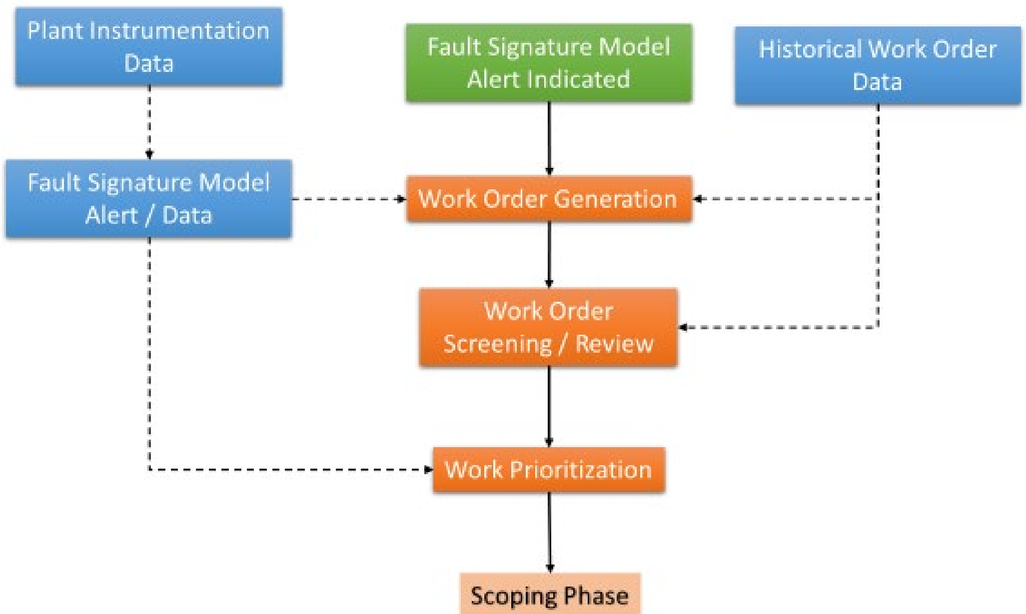

Figure 57 - Screening Phase Data Flow

The NDP integrates the work screening phase into the digital application by applying the data that would be used to generate a work notification based on the associated fault. This includes the date, time, diagnosis, and recommended corrective actions. Instead of performing separate approval of the work notification, the NDP automatically categorizes faults identified by the model as valid based on the high accuracy of AI/ML modeling techniques. This design requires that the CBM tools identify issues with high accuracy, consider design requirement conservatism, and have sufficient instrumentation to support the intended conclusions.

A WO created at a utility would typically be evaluated by stakeholder organizations to determine the work classification, priority, and impact to the plant. This can be implemented flexibly by the Automated Work Management Process by allowing for reviews on behalf of the various stakeholder organizations. Once the reviews are completed, the NDP output can be printed as a report and attached to a plant-generated WO or used to generate an automated WO within the utility's Enterprise Management System (EMS). It is recommended that integration into the EMS be considered long-term in order to increase process efficiency and reduce the duplication of work. However, using a digital platform as a hub for reviewing failure alerts and reporting, without integration to an EMS, still improves plant identification of degradation risks and improves the data quality used to support generation of a work order.

In addition, data such as work order history, component design and licensing requirements, and material estimates can further support the understanding of the work to be performed and the impact to the plant. This data should be strongly considered for inclusion in any digital platform as doing so can support stakeholder reviews, scheduling, supply chain services, and prioritization. Where possible, it is recommended that digital platforms emulate processes performed by plant stakeholders to maintain integrity with the existing plant processes and streamline integration into existing supporting services.

During development of the digital platform, it was understood that addressing all severities of plant conditions immediately would alter the scope and design intended for this project. Emergent work performed to maintain plant

Operability requires a different level of rigor when compared to evaluating long-term degradation risks. The scope of this project was to utilize automated work management as a tool to enhance the maintenance strategy for Salem's CWS. A platform designed to address emergent or high priority work needs to consider data refresh rates, plant stakeholder accessibility, architecture redundancy, and application availability in a different manner than one used as a tool to address lower priority work streams. It is recommended that digital platforms be designed and implemented using a phased approach to focus on the highest value services first, with consideration of scalability or expansion once the initial services are functional.

As a part of this project, the work recommendations such as historical WO's or maintenance plans made by the software are manually mapped using personnel who have Plant Engineering and Operations experience as reference. This is considered a best practice for small-scale implementation where limited operating scenarios are considered and accuracy within the limited subsets of events provides the greatest value. For large-scale implementation of an Automated Work Management Process, ML and NLP techniques should be designed and implemented to minimize manual feedback required to provide accurate recommendations. The ML/NLP techniques add significant value when user inputs are incorporated as a feedback loop for model improvement, and work recommendations are ranked based on the scope of a given anomalous condition.

One incremental goal within the scope of this project, was to utilize input from a CBM model to identify the need to perform an existing maintenance plan. The maintenance plan is selected to resolve the identified component fault by way of mapping fault signature model outputs to the corresponding maintenance plan. The goal of linking a fault signature model outputs to existing maintenance plans was chosen because it requires fewer interfaces with stakeholder organizations and is easier to prioritize and schedule as compared to WO's that are not tied to a maintenance plan. This approach also provides a template for how time-based maintenance can be deferred to on-condition maintenance which is another goal for the project.

Based on the above, it is important to understand the types of WO's which cannot be fully addressed through template maintenance plans due to process restrictions. Plant procedures identify stakeholder reviews by SME's when the subject of the WO has unique requirements. These reviews identify additional work steps or contingencies that are required to align with plant commitments and regulations. Table 10 is a partial list of scenarios or topics that may require review by specialty SME's.

Specialty Subject Review Topics		
Environmental Review	Boric Acid Leak	Radiological Exposure Risk
Impacts to Plant Design Basis (Modifications, Setpoint Changes)	Shutdown Safety System	Industrial Safety Review
FLEX (Fukushima Mitigation) Related	Foreign Material Exclusion (FME) Related?	Regulatory Consequences
Emergency Response	Reactivity Affected?	Security Related?
Fire Watch Required	Potential Impact to Fuel Reliability?	

Table 10 - Example Unique Subject Matter Review Topics

Due to the large number of potential specialty review areas that needed to be considered, the preferred approach for an initial implementation of an Automated Work Management Process would be to avoid components located in systems which impact multiple specialty subject matter topics. As noted within this report, the focus of this project was the CWS at Salem Nuclear Power Plant. The CWS is a non-safety related system with minimal links to the specialty subject matter topics above. As a result, the corresponding reviews will be less complex and require fewer stakeholder organization interactions. It is recommended that implementation of an automated process including complex systems with special considerations be integrated with the relevant EMS data to either automatically screen events against unique scenarios or to force reviews by assigned SME's.

As noted above, the Fault Signature Model needs to provide an estimate of when work must be performed to allow for coordination with an overall work management process. One method for performing this assignment is using a matrix-based scoring method to classify the urgency of the work. This method was selected to simplify scheduling decisions as much as possible based on the CBM model output and the underlying component data from the plant's EMS. The outputs of the Fault Signature Model and component classification data from a utility's EMS provides the column classifications for prioritization. The row classifications are based on plant programs and conditions that identify the criticality of a component in regard to maintaining plant operability and adherence to codes and standards.

Prioritization Matrix

		Group 1			Group 2		
		Column 1	Column 2	Column 3	Column 4	Column 5	Column 6
		Not Functional	Not Functional	Not Functional or Functional	Functional	Functional or OM	Functional or OM
		TS Inoperable	TS Inoperable	TS Inoperable or Non TS	Non TS	Non TS	Non TS
		Shut down actions governed by TS	Compensatory (Comp.) actions required to maintain TS or PRA available	Comp actions required to maintain PRA available	Monitoring or Comp Actions required to maintain PRA Available	PRA Applicable No actions required to maintain PRA available	PRA Availability not applicable
Work Significance	Comments	Corrective	Corrective	Corrective or Deficient	Deficient	Deficient or Other	Deficient or Other
Immediate Threat to Public Health and Safety		1	2	3	4	5	5
Technical Specification (TS) System/Component		1	2	3	4	4	5
Risk Significant (Plant Trips, SPV's, etc.)		2	2	3	4	4	5
Operations Workaround (Comp Measure)		2	3	4	4	5	5
Accelerated Degradation		2	3	4	4	5	5
Tech Spec/ORM/ODCM (> 14 Days to Req'd Action)		3	3	4	4	5	5
Significant Economic Risk (Outage Risk, Critical Spares)		3	4	4	4	5	5
Quality-Related (Including Regulatory Commitments)		4	4	4	4	5	5
Burden to an Organization (Impacts to EOP's, AOP's)		4	4	4	5	5	5
BOP Reliability/Nonregulatory Programs/Nonplant SSC		5	5	5	5	5	5

Figure 58 - Work Prioritization Matrix

Key	
Priority Score	Target Window/Next Action
1	Won't be handled as predictive. Already at an alarm state. Immediate Corrective action per existing site process.
2	Won't be handled as predictive. Already at an alarm state. Immediate Corrective action per existing site process.
3	Could be handled as predictive, in cases where the equipment could be judged as "Functional". If equipment is "Non-Functional", Immediate Corrective action per existing site process. A suggested date within 90 days with link back to Fault Signature Prognosis.
4	Work within the next appropriate work week for the Functional Equipment Group. Assume no more than 180 days, but Fault Signature Prognosis may allow for this to be further pushed out. Attempt to target Functional Equipment Group (FEG) window, if suitable.
5	Needs to be in a future Functional Equipment Group (FEG) Outage week, or as fill in work. Can be easily autoassigned a date, or put on permanent trend to possibly not even be part of an outage work.

Figure 59 - Work Prioritization Matrix Key

Figure 54 and Figure 55 identify the criteria for consideration and the associated priority score that has been assigned for each category. As noted above, at this time the Automated Work Process does not intend to resolve plant component failures requiring immediate corrective action. Work requiring immediate action will be handled through existing site processes for emergent work until greater integration with the Automated Work Management Process occurs. Based on that logic, items in Group 1 (Columns 1 through 3) will be outside the Automated Work Management Process except for degradation events which fall in column 3. Column 3 is unique due to compensatory measures required for Probabilistic Risk Assessment (PRA) program needs; however, degradation can still be addressed in the automated system by interfacing with the appropriate stakeholder organizations.

Work that is in scope for the initial phase of the Automated Work Management Process falls in Group 2 (Columns 4 through 6). The targets for this type of work are deficient conditions which have minimal compensatory actions and do not impact plant or component operability. For these conditions, we will use the lowest priority score available based on component data to determine the approximate timeline in which the work must be performed, based on the component function and requirements. Each target score will have a target work window, which provides a baseline expectation of how long work can be deferred. This target work window can be overridden by the prognosis output of the CBM model.

In summary, the Automated Work Management screening phase is critical for linking the output of a CBM model to the beginning of WO generation. The considerations documented above are critical for understanding the work required, the deficient plant condition, and the necessary corrective measures. These considerations mirror those performed by a plant or utility within their existing work management processes; thus, it is anticipated that the screening can be performed in parallel procedural paths to the existing process.

6.3 Scoping & Planning Phase

The next two phases of the Work Management Process are the scoping and planning phase. In the scoping phase, plant stakeholders refine the information proposed during the screening phase, aligning any initial estimates provided during screening to better defined material and labor commitments. In the planning phase, the work defined during the scoping phase is reviewed to ensure accuracy and completeness. Work Planning and Work Management personnel are heavily involved in these phases to verify that the work can be performed using the outputs of the screening phase. A brief list of considerations for the scoping and planning phases are listed below.

- Validation of Labor Resources
- Stakeholder Review of Proposed Scope
- Identification/Ordering of Materials
- Identification and Resolution of Restraints to Support Work Planning

In an Automated Work Management Process, the scoping phase focuses on refining the automated recommendations generated during the screening phase. The goal of this phase is to ensure that the work is appropriately scoped and stakeholder input is fully incorporated. As noted above, the Automated Work Management Process should include automation of the stakeholder input to the greatest extent possible to eliminate manual interaction with the process. In the scope of this project, manual input is used to validate model outputs and provide confidence that the model is providing appropriate outputs. Figure 60 shows the simplified process flow of the scoping phase.

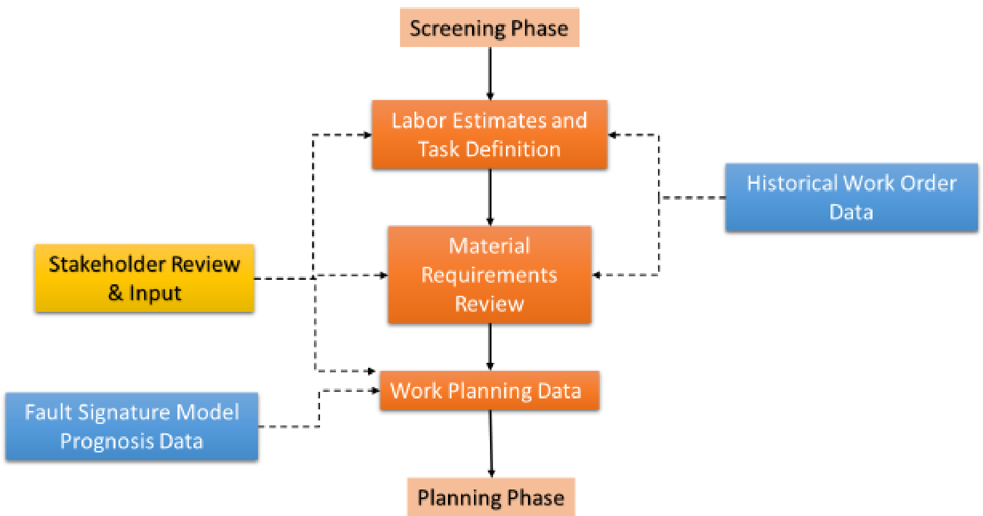

Figure 60 - Automated Scoping Phase Data Flow

The goal of the planning phase is to verify the scoping already performed, and prepare work instructions for use by craft personnel, which requires aligning maintenance plan with the corresponding instructions and work procedures. These instructions provide the template to perform the intended work and can be used by craft personnel to perform pre-job walkdowns and confirm the scope of work to be performed. Figure 61 shows the process flow within the planning phase of the Automated Work Management Process.

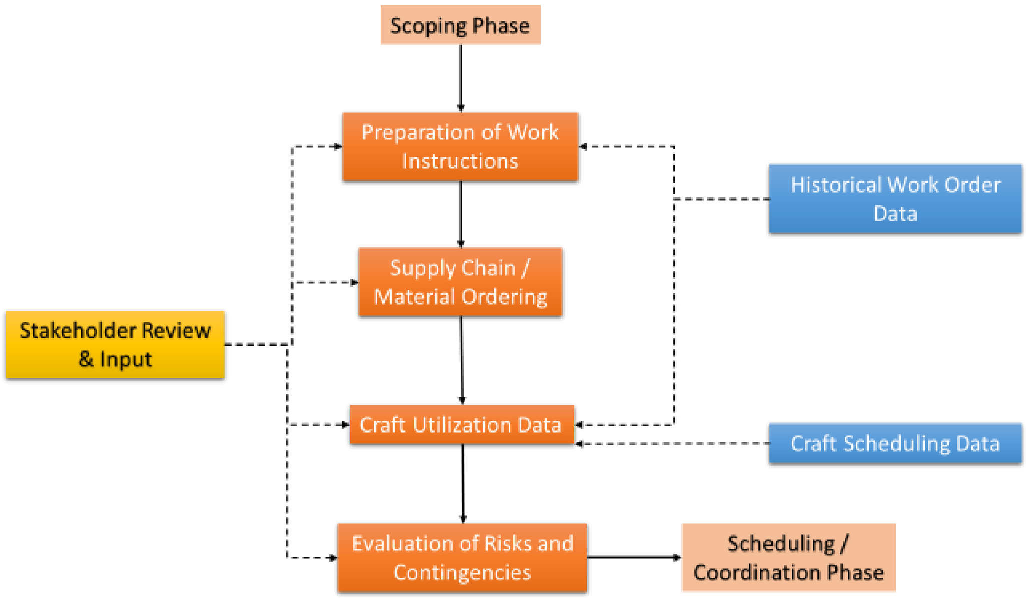

Figure 61 - Automated Planning Phase Data Flow

After receiving an indication of a fault from the CBM models, the first reviewable output of the Automated Work Management Process is the extent of the condition. Data outputs from CBM tools will provide background on the identified fault as well as information on the component and sensors that led to the alert. Once the suspected fault is identified, the next step is to confirm the maintenance solution required to resolve the fault. Corrective maintenance could include performance of an existing preventive maintenance plan or historical WO. As noted above, using historical maintenance as a recommendation for correcting an identified fault simplifies the work planning process within the digital platform. Specifically, using a template WO or maintenance plan provides the system data to support the generation of work instructions, identify the craft personnel resources needed, and ascertain the necessary stock to perform the work.

Once the template work order or maintenance plan is recommended by the NDP, plant personnel can begin reviewing and modifying the task to suit the material condition of the component. The review process includes screening the CBM model outputs (e.g., diagnosis and sensor data) for scheduling and scoping, thus ensuring that the NDP recommendation is reasonable. This review also refines the labor estimates for craft personnel and the task definition for instructions to ensure that the work plan is suitable, while still minimizing the work planning time by utilizing historical template WO's or maintenance plans. In addition, the manual review allows site stakeholders to confirm the labor, materials, and instructions required to perform the work. It is recommended that review changes be identified and tracked within the platform for use in a feedback loop to improve the model's decision making using ML or NLP, as appropriate.

The review of materials used significantly impacts the planning and scheduling associated with PM tasks because parts must be ordered, stocked, and maintained for use within the plant. As a result, supply chain considerations are critical for working efficiently within an Automated Work Management Process. Thus, PKMJ reviewed CWS WO's for the pump and motor sets across several PM tasks to determine the data quality currently available. The review focused on identifying the required and contingent parts for a given task type over the last two performances of the task per pump and motor set (maximum of twenty-four (24) work orders). The team's review assumed that "Required" parts were used across multiple WO's corresponding to a particular task type and that "Contingent" parts were used only once per task type. This simple assumption was used to benefit a manual review with the knowledge that data processing could provide insights into appropriate thresholds. The output of the manual review is shown in Table 11.

Maintenance Plan	# of Work Orders	# of Unique Parts	# of Required Parts	# of Contingent Parts
18M CW PUMP UNDERWATER INSP	24	42	20	22
6M LUBE SAMPLE: ALL CW PUMP MOTORS	4	2	2	0
54M LO FLOW SW: CAL SWITCH & FUNCT CHK	21	5	4	1
4Y CAL / CW PUMP PRESS SW	24	9	4	5
3Y INSP CIRC WATER PUMP MTR	22	21	9	12
6Y CWP MTR REPLACE MOTOR	24	191	39	152
6Y CIRCULATOR PUMP CHANGE OUT	24	211	59	152

Table 11 - Parts Statistics per Maintenance Plan

The results of the review regarding parts consumed during PM tasks above provide insight into how the parts consumed for work are consistent for smaller maintenance tasks, whereas more contingencies are required for larger tasks. In particular, data shows that a large amount of contingent parts are required for the pump and motor replacement tasks. WO material predictions for instrumentation and oil-related tasks can be expected to closely aligned thanks to these insights. In addition, instances of consumables (markers, adhesive) and tools (gloves, ladders) in these WO's create a challenge for work-specific analytics. Consumables and tools not mission-critical to the work being performed can obscure insights into common components required for work (i.e., searching for all work that requires markers has limited analytic value). However, consumables such as oil are valuable to the analytics of the oil sampling and inspection tasks. It is recommended that non-critical consumables and tools be omitted from analyses to improve the accuracy of the results.

Another important insight gleaned from the manual review was the identified utilization of replacement or equivalent components in WO's over time. Several instances of consuming new parts of the same type (e.g., concrete anchors, valves, bolts, etc.) were identified. In various cases, the WO's required both the old and new parts to be retrieved from the warehouse to the work location – an inefficiency addressable by an automated process. It was also identified that a limited amount of Safety-Related stock was used interchangeably with Non-Safety Related stock in Non-Safety applications. This is an occasional practice in supply chain organizations when a surplus of Safety-Related stock is available for use, or if an emergent issue requires immediate action and the Non-Safety stock has a long lead-time. Safety-Related stock has a higher per unit cost than equivalent Non-Safety components, due to underlying quality requirements. The automated system could help relieve supply chain lead-time requirements by identifying upcoming maintenance work sooner, thus saving plants money by allowing them to purchase and use Non-Safety Related stock.

6.4 Scheduling / Coordination Phase

Scheduling and coordination of work under a Work Management Process is critical for ensuring that stakeholder organizations are prepared to perform work at the assigned time. This phase of the process requires input from organizations such as Operations, Radiation Protection, Security, Supply Chain, Work Management, and others. These organizations work together to assign resources to support the work and verify that the required materials will be available at work start. This is performed with insight into the plant's status so that consideration of work windows and potential outages can lead to increased efficiencies within various work scopes.

In an Automated Work Management Process, the goal is to automate as many stakeholder organization resource decisions as possible to support large-scale implementation of the automated system. Upon review of the planning phase's output related to labor and materials, an understanding of the necessary resource commitments is achievable. These considerations can be utilized to their full effect when simultaneously managing work scheduling for several stakeholder organizations – a task made difficult for humans thanks to cross-organizational dependencies and competing priorities. Figure 62 shows the relationships among these various steps.

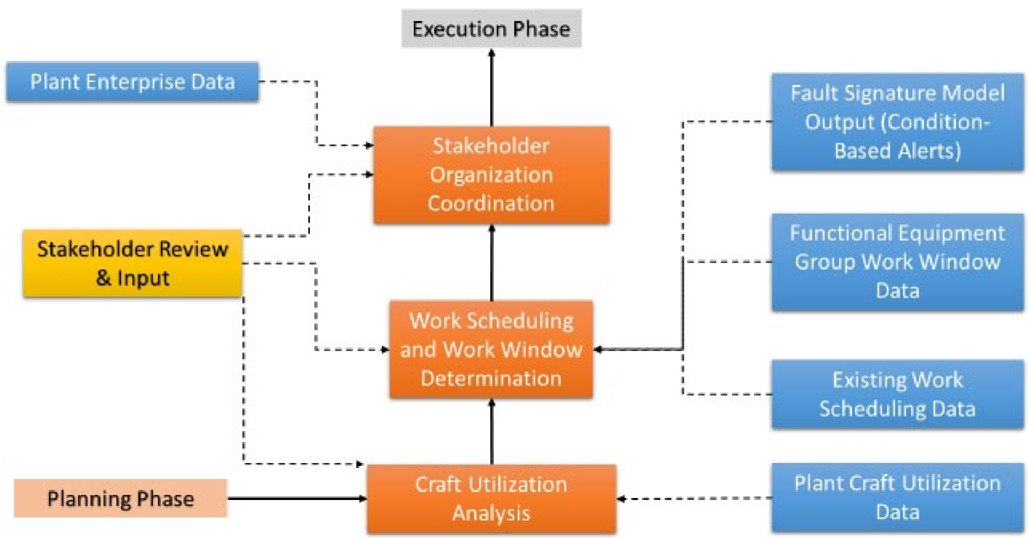

Figure 62 - Automated Scheduling and Coordination Phase Data Flow

Scheduling heavily depends on data from a utility to define the utilization of their craft personnel for planning work. Due to challenges associated with system access, this project did not obtain craft data for PSEG. However, several sharable design recommendations were identified during development of the process.

PKMJ worked with PSEG to understand their maintenance windows through the assignment of Functional Equipment Groups (FEGs). These maintenance windows, known at the utility as "FEG windows", represent the planned times when equipment would be removed from service for maintenance, calibration, and repair. Where possible, failure risks should be identified with sufficient lead time to schedule corrective work within a FEG window. This would foster efficiencies for craft personnel by enabling the utilization of shared resources. In addition, Operations personnel could better control impacts to the plant if work was clustered together to support the establishment of a common clearance. Scheduling work should be aligned based on the failure risk severity, component performance data, and component requirements to establish the appropriate timing to perform the work compared to other planned maintenance.

Another recommendation is to integrate, to the extent possible, craft utilization data into the work scheduling process. This would allow a better understanding of which tasks require an influx of labor personnel, as well as better management of the required skillsets. A maintenance program reliant on CBM must consider the altered labor profile, as work will be performed less frequently or on variable frequencies, depending on the component condition. The consistency of skill-specific work tasks is expected to be a challenge for managing a CBM program. Prognosis outputs from CBM tools provide insights into planning work within a CBM program, but the specific handling of this issue must be adjusted based on the extent of the maintenance scheduling variance due to online monitoring of plant components.

Once the work is fully scheduled and the labor resources defined, the process must verify that stakeholder organizations such as Chemistry, Security, and Operations are notified of the work and can support it. These organizations will have opportunities to review the WO in the previous phases; however, they will need to formally assign resources to the work. For example, Chemistry will need to address ALARA concerns for radiation protection, Security may need to escort vendors into the plant to support maintenance, and Operations will need to post clearances for the planned work. Each organization must be notified of their respective requirements based on the WO, review the output of the automated system, and act accordingly. One area for future enhancement is the use of NLP to improve the digital platform's recommendations in these areas.

The scheduling and coordination phase of the Automated Work Management Process is one of the best opportunities for utilizing tools to streamline the work process. Coordination among several different stakeholders is often difficult, and an automated process could minimize conflicts and challenges. Scheduling for work in a complex work environment represents an opportunity for automated processing to replace manual decision making, which is often based on procedural considerations. In an initial implementation, establishing the scheduling so as to fit within resource loading is a challenge due to data availability. Once the data challenges are resolved, gradual expansion of automated scheduling to encompass more situations can be considered.

6.5 Execution and Post-Work Analysis Phases

The work execution phase is the phase in which work is performed in the field, whereas the post-work analysis phase is that in which performance data and feedback are captured. Today, these phases are performed without automated intervention, but certain aspects of work execution should be tracked via a Automated Work Management Process. When field conditions deviate from those assumed in the scoping, planning, or scheduling phases, it is imperative that feedback be provided to the various AI/ML algorithms to improve future performance. Figure 63 shows the conclusion of the process flow related to the Automated Work Management Process.

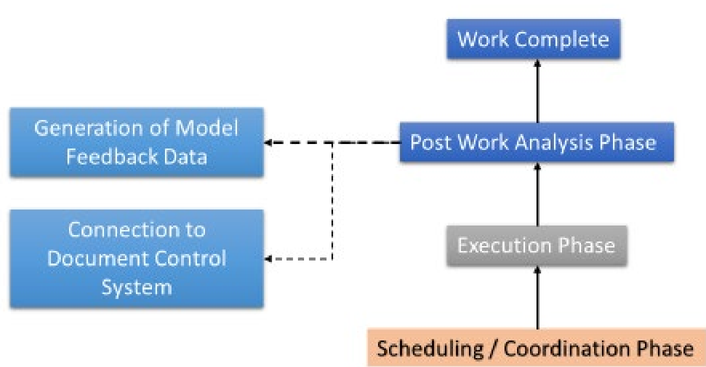

Figure 63 - Work Execution and Post Work Analysis Phases

Recommendations for tracking performed during the execution phase include enhanced tagging for Measuring and Monitoring Equipment (MME) and tools being integrated into the digital platform, providing a record of when calibrated instruments are used. In addition, this could be combined with other plant digitalization efforts (e.g., RFID tagging of tools) to provide real-time data on how long work takes to perform in the field, thus supporting craft personnel utilization estimates within a digital platform. Detailed work performance data not historically collected, due to the burden of documenting different datasets, could become a powerful tool for supporting efficiencies gained when implementing an automated process.

A feedback mechanism is also recommended for incorporation into the WO closeout process. Currently, work organizations close out paperwork by entering the scanned WO – with completion notes and logs – into a document repository. These completed WO's should also include input fields that can give the automated system recommendations and work performance feedback. The following fields are examples of possible post-work inputs needing addressed:

- Scope of Work Performed
- Material Condition
    - Determine if this aligns with the alarm outputs
- Correct Materials
    - If not, the list of changed materials should be inputted
- Work Duration
- Work Instructions Acceptable
    - Detail any changes including pen-and-ink markups
- Other Feedback or Notes for Future Repetitive Tasks

The initial implementation of an automated process will require procedural changes to enforce the feedback structure. Instead of via free text, the feedback will be performed using distinct selections or valid values to provide structure for model feedback and automation. Over time, the process is expected to improve in accuracy by learning.

Aspects of the Automated Work Management System will be common across all work scopes, further improving the rate at which the system learns from experience. For those scopes unique to a particular component or component type, opportunities exist to incorporate analyses and best practices from industry groups or manufacturers in order to speed up the model's maturation process. Without experimental data, learning can still be performed on the available data set in order to reach a conclusion at a relative level of confidence. With the implementation of multiple CBM models within the same platform, there is confidence that different statistical and analytical methods can provide accurate predictions of component conditions.

6.6 Summary

Automation of the Work Management Process is an area for improvement in the nuclear industry. In many cases, existing processes that support work management are compartmentalized and would benefit from digitalization and the removal of manual administration. With automation, the ability to store and analyze data becomes more relevant and supports the implementation of AI/ML tools to enhance insights retrievable from the data. Large data sets such as those for WO and stock data are difficult to conceptualize during manual review; however, available tools to automate these reviews can be incorporated into a process that benefits the industry. The goal of the overall Automated Work Management Process is to reduce the burden on NPP's when work must be performed and the process described in this report lays the foundation for that improved process.

7. SUMMARY AND CONCLUSION

This report presents the R&D performed by PKMJ in collaboration with INL and PSEG to develop a centralized digital platform. The digital platform was developed to provide the structure for an industry solution for performing data-driven services supporting improved equipment reliability and PMO. An industry platform offers new, innovative solutions and efficiencies that enhance existing nuclear processes and reduces overall O&M costs.

Work under this effort began with the implementation of a risk-informed approach for optimizing component maintenance strategies for the CWS at Salem. This was done by combining enterprise-specific data with industry data to enable SME's to make informed decisions on the overall maintenance strategy for the CWS pumps and motors. The NDP quickly provides value to utilities by allowing PMO reviews to be performed digitally with visualizations that support the conclusions readily available. Salem will realize this value through $4.37M in CWS-maintenance-related savings over the next six years.

The risk-informed approach is further enhanced by the CBM and APR modeling developed to further support equipment reliability for Salem's CWS. This project examined the framework necessary to identify degradation risks for components prior to impacting operating plant components. As a result of this effort, PSEG enhanced their monitoring capability for an important Balance of Plant (BOP) system and improved their ability to predict and diagnose failure events. The capability to mitigate degradation risks before they become serious events, schedule corrective maintenance, and perform the supply chain services required for the maintenance, is a major achievement of this research.

PKMJ's NDP provides the structure for centralizing and digitalizing processes within an NPP. For example, Engineering, Operations, and Maintenance personnel can review, whether individually or simultaneously, component degradation risks to determine how best to address those risks. This project describes how to create a digital platform that is similar to the PKMJ NDP and allows for expansion across systems or fleets to provide an optimized approach to various processes. If a utility or fleet has already developed their own processes for evaluating equipment reliability, those tools could be integrated into the platform to supplement the insights received.

As degradation risks are identified and prioritized for components, the automation of work notifications or WO generation is critical for simplifying the maintenance process. Repetitive maintenance tasks place a significant administrative burden on organizations that operate nuclear facilities. Automation of simple decisions within the process reduces the burden on the personnel who must manage those processes. Utilities can implement the automated processes described in relation to this project in order to optimize maintenance programs while still ensuring their effectiveness at maintaining healthy plant conditions.

As this research project comes to an end of its period of performance, the activities performed within its scope lay the foundation to expand the technologies and capabilities developed to other processes, such that the nuclear industry can further benefit from the work performed. This research provided direct value to the utility partner (PSEG) and enabled in-depth research into valuable solutions that are expandable to additional components, sites, and fleets. Collaboration among industry stakeholders under this project uncovered many different opportunities for future improvements. Partnership between utilities, research organizations, and vendors will be critical for nuclear as the industry moves towards plant digitalization via efforts such the effort described herein.

8. REFERENCES

[1] V. Agarwal, *Risk-Informed Condition-Based Maintenance Strategy: Research and Development Plan,* Idaho Falls: Idaho National Laboratory, INL/LTD-18-51448, 2018.

[2] PKMJ Technical Services, LLC, Idaho National Laboratory, and Public Service Enterprise Group (PSEG) Nuclear, LLC, "Integrated Risk-Informed Condition Based Maintenance Capability and Automated Platform: Technical Report 1 (PKM-DOC-20-0013)," Department of Energy, 2020.

[3] Microsoft, "SLA Summary for Azure Services," June 2021. [Online]. Available: https://azure.microsoft.com/en-us/support/legal/sla/summary/. [Accessed 30 June 2021].

[4] Microsoft, "Grant limited access to Azure Storage resources using shared access signatures (SAS)," 28 December 2020. [Online]. Available: https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview#controlling-a-sas-with-a-stored-access-policy. [Accessed 29 June 2021].

[5] Microsoft, "How to Manage SAS Token Which Has Already Been Shared Out?," 17 December 2020. [Online]. Available: https://docs.microsoft.com/en-us/answers/questions/201900/how-to-manage-sas-token-which-has-already-been-sha.html. [Accessed 29 June 2021].

[6] Microsoft, "Security Recommendations for Blob Storage," 1 March 2021. [Online]. Available: https://docs.microsoft.com/en-us/azure/storage/blobs/security-recommendations. [Accessed 29 June 2021].

[7] Microsoft, "Simplify Your Lakehouse Architecture with Azure Databricks, Delta Lake, and Azure Data Lake Storage," 6 January 2021. [Online]. Available: https://techcommunity.microsoft.com/t5/analytics-on-azure/simplify-your-lakehouse-architecture-with-azure-databricks-delta/ba-p/2027272. [Accessed 29 June 2021].

[8] Microsoft, "Serverless Web Application," 2021. [Online]. Available: https://docs.microsoft.com/en-us/azure/architecture/reference-architectures/serverless/web-app. [Accessed 29 June 2021].

[9] Chen T, Guestrin C., "XGBoost: A Scalable Tree Boosting System," in *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.

[10] P. Harrington, Machine Learning in Action, Simon and Schuster, 2012.

[11] Lundberg, Scott, and Su-In Lee, "A Unified Approach to Interpreting Model Predictions," *arXiv Preprint arXiv:1705.07874,* 2017.

[12] PSEG, *Configuration Baseline Document for Circulating Water System, DE-CB.CW-0028(Z), Rev. 0, 1-1.*

[13] Idaho National Laboratory, PKMJ Technical Services, LLC, and Public Service Enterprise Group (PSEG) Nuclear, LLC, ""Machine Learning and Economic Models to Enable Risk-Informed Condition Based Maintenance of a Nuclear Plant Asset (INL/EXT-21-61984)," Department of Energy, 2021.

[14] W. Feller, An Introduction to Probability Theory and Its Applications, Volume 1, 3rd Edition, Wiley, 1968.

[15] W. L. Winston, Operations Research: Applications and Algorithms, 2nd Ed., Boston: PWS-Kent Publishing, 1991.

[16] Harry F. Martz and Ray A. Waller, Bayesian Reliability Analysis, New York: Wiley, 1982.

[17] L. Kleinrock, Queueing Systems, Vol I: Theory, Wiley, 1975.

[18] R. Nelson, Probability, Stochastic Processes, and Queueing Theory, 1995: Springer-Verlag.

[19] Institute of Nuclear Power Operations (INPO), "AP-928, Online Work Management Process Description," Rev. 5, September 2017.