Network Design in MRO Inventory for Multinational Conglomerate

by

Elie Maroun

and

Gavin Calarco

Submitted on May 8, 2026 in Partial Fulfillment of the
Requirements for the Degree of Master of Applied Science in Supply Chain Management

at the Massachusetts Institute of Technology

ABSTRACT

This project addresses the inefficiencies of a decentralized Maintenance, Repair, and Operations (MRO) inventory network for a large, vertically integrated conglomerate operating across Canada and the United States. Driven by a fragmented stocking strategy, the organization's existing MRO supply chain suffers from duplicate inventory, inflated safety stock levels, and limited visibility into total material and freight costs. To evaluate the trade-offs between cost, service reliability, and downtime risk, this study pairs a multi-dimensional Stock Keeping Unit (SKU) segmentation framework with a Mixed Integer Linear Programming (MILP) network optimization model. By utilizing K-Means clustering across ABC (spend), XYZ (volatility), and VED (criticality) dimensions, the model successfully isolated 241 vital, high-spend SKUs as prime candidates for centralization. The MILP simulation demonstrated that transitioning from 15 independent facilities to a hybrid centralized architecture anchored by two hubs (one in Canada and one in the U.S.) yields an optimal balance of cost and operational resilience. By leveraging the square root law of risk pooling, this two-hub configuration reduces required safety stock by 33% in Canada and 12% in the United States, decreasing the base total landed cost from $211,517 to $140,851 and generating $70,666 in annual structural savings. When inventory pooling savings ($90,473/yr) and procurement consolidation benefits ($14,829/yr) are included, the total extended annual saving reaches approximately $175,968 – $164,389 within Canada, and $11,578 within the United States. This is paired with a one-time capital release of $361,894 from rationalizing distributed on-hand inventory – $346,815 within Canada, and $15,079 within the United States. To sustain these improvements, the project recommends institutionalizing automated SKU deduplication, shifting to Free Carrier (FCA) procurement terms to leverage the company’s transportation fleet for regional short-haul multi-destination routes, and adopting the targeted hub redesign.

Capstone Advisor: Dr. Ilya Jackson
Title: Research Scientist, MIT Center for Transportation & Logistics

TABLE OF CONTENTS

1 INTRODUCTION 3
1.1 Motivation 3
1.2 Problem Statement and Key Questions 3
1.3 Project Goals and Expected Outcomes 4
2 STATE OF THE PRACTICE 5
2.1 MRO Supply Chain Challenges 5
2.2 Demand Forecasting and Segmentation 5
2.3 Network Design and Centralization 6
2.4 Total Landed Cost Analysis 7
2.5 Relevance to Problem 7
3 METHODOLOGY 8
3.1 Methodology Selection 8
3.2 Data Collection and Preparation 8
3.3 Application of Methodology 9
3.3.1 Multi-Dimensional SKU Segmentation 9
3.3.2 MILP Model Design 9
3.3.3 Network Routing and Internal Logistics Simulation 10
3.3.4 Scenario Design 10
4 RESULTS AND DISCUSSION 11
4.1 Current State and Segmentation Results 11
4.2 Network Optimization Findings 12
4.3 Inventory and Cost Implications 14
4.4 Sensitivity Analysis: Hub Robustness and Designation Costs 15
4.5 Predictive Modeling Performance and Demand Characterization 16
4.6 Inventory Policy Dynamics: Risk Pooling and the Square Root Law 16
4.7 Limitations 17
5 RECOMMENDATIONS 18
5.1 Management Recommendations 18
5.1.1 Immediate Tactical Actions: Master Data and Segmentation 18
5.1.2 Strategic Network Redesign 18
5.1.3 Leveraging Vertical Integration for Internal Logistics 18
5.2 Future Work 18
REFERENCES 20

1 INTRODUCTION

Background

The project sponsor is a large, vertically integrated conglomerate headquartered in Canada with extensive operations across Canada and the United States. Established over 50 years ago, the organization operates a diversified portfolio of businesses. This project focuses on operations within the sponsor's production and logistics divisions. (Anonymous, personal communication, October 15, 2025).

1.1 Motivation

The sponsor team is motivated by the need to improve operational efficiency, increase cost transparency, and develop a scalable Maintenance, Repair, and Operations (MRO) inventory supply model that supports long-term growth. As operations expand across multiple facilities in close geographic proximity, MRO activities remain highly decentralized, leading to duplicate inventory, inconsistent master data, and limited visibility of total material and freight costs. As a result, inventory levels are higher than necessary, total costs are difficult to quantify, and opportunities for consolidation remain unrealized. These issues create inefficiencies, tie up working capital, and limit the organization's ability to fully understand true landed costs or leverage economies of scale.

Through this project, the sponsor seeks to assess whether a more centralized MRO inventory structure can maintain high service reliability while reducing inventory and transportation costs.

1.2 Problem Statement and Key Questions

To address these issues, the project will assess the potential benefits and trade-offs of centralizing MRO materials and redesigning the physical material flows within the network. By simulating centralized and decentralized stocking scenarios, the project will evaluate trade-offs among cost, service level, and downtime risk. In this context, a machine within a factory is in downtime if it is waiting on a replacement part to be installed before it can run at its expected output. Downtime risk refers to the risk associated with limiting the quantity and/or diversity of replacement parts available at a facility.

The findings will help the sponsor optimize inventory management, enhance supply chain visibility, and build a more resilient, efficient operating model. The goal is to determine whether a more coordinated model could reduce inventory and freight costs while maintaining high service reliability and minimizing downtime risk.

In this context, the key questions are:

1. How can the sponsor organization redesign its MRO inventory structure to reduce redundancy and improve visibility across sites?
2. What level of centralization provides the best balance between cost, service level, and operational resilience?
3. How can the sponsor leverage its logistical assets (trucks, dispatchers, etc.) to reduce freight and total landed costs?

To explore these questions, the project will proceed through several stages. To begin, the current MRO process will be mapped to understand how parts are sourced, stored, and managed at each facility. Data from the Enterprise Resource Planner (ERP) system will then be collected and reviewed for consistency and completeness. Using this data, models will be developed to examine: (1) centralized and decentralized stocking scenarios, (2) utilization of internal freight, and (3) key performance indicator implications. The

results will be validated with sponsor stakeholders, and practical recommendations will be developed to support more efficient inventory management, improved visibility, and long-term scalability.

1.3 Project Goals and Expected Outcomes

By comparing centralized and decentralized stocking structures as well as internal freight utilization, the project will determine how to reduce duplicate inventory and enhance transparency of total material and freight costs while maintaining service reliability.

The deliverables to the sponsor will include:

1. A suite of pre-run scenarios evaluating centralized and decentralized MRO scenarios, showing cost, service level, and downtime trade-offs.
2. Recommendations for optimizing inventory management and improving network coordination.

2 STATE OF THE PRACTICE

The primary goal of this capstone is to determine how the sponsor company can reduce costs within its MRO supply chain while maintaining target service levels. To better understand the theory behind this and related challenges, we reviewed a few key areas of MRO supply chain academic literature: (1) current and historic challenges, (2) demand forecasting and segmentation methods, (3) warehouse network design and the role of inventory centralization, and (4) total landed cost analysis.

2.1 MRO Supply Chain Challenges

MRO items are goods required to maintain a facility's production output but do not become part of the finished product. These include consumables such as machine lubricants and industrial components, such as bearings or belts. Determining demand requirements of MRO items is difficult due to four main complexity drivers: "the high number of parts managed; the presence of intermittent or lumpy demand patterns; the high responsiveness required due to downtime cost for customers; and the risk of stock obsolescence" (Bacchetti & Saccani, 2012).

As their demand characteristics are anything but traditional, typical demand forecasting methods cannot be relied upon for MRO parts. The number of distinct MRO Stock Keeping Units (SKUs) managed within companies that are comparable to our sponsor in size and form is typically in the tens of thousands (Claure & Ovalle, 2023). This makes it infeasible to optimize the supply chain for every individual part. These facts, along with managers' aversion to MRO-part stockouts that drive asset downtime, mean that the capital dedicated to holding MRO inventory is immense.

2.2 Demand Forecasting and Segmentation

Item forecasts are made to advise future ordering decisions, because there is a non-zero lead time associated with the order-to-delivery cycle. Demke et al. (2025) found limited research on forecasting methods tailored to the unique characteristics of MRO parts. Traditional time-series methods, such as Moving Average and Exponential Smoothing, are sometimes used because of their familiarity and widespread use in conventional supply chains. However, they fail in MRO supply chains, as they were developed for items with steady or smooth demand. These methods overestimate demand, leaving companies that utilize them with excess MRO parts held in inventory (Demke et al., 2023).

Legacy methods for intermittent demand, such as Croston's method, are currently used in practice. Willemain et al. (1994) show that Croston's method is significantly superior to exponential smoothing, as it forecasts with higher accuracy, lower variance and error, and greater overall robustness against underlying assumptions. Supply chains have evolved since Croston's method was developed and matured. Demke et al. (2025) argue that univariate methods fail to consider the dimensionality and complexity of MRO-parts demand. Univariate forecasting methods predict future values of a single time-dependent variable based solely on its own historical data, rather than including other related factors. They propose developing and maturing machine learning models that leverage vast amounts of available MRO data to inform strategic supply chain decisions, such as optimal warehouse locations and dynamic inventory policies.

The vast number of SKUs within MRO supply chains makes it infeasible to make effective decisions at the SKU level. To approach decision-making for these parts, SKUs can be segmented into groups, allowing group-wide decisions. Segmenting parts allows supply chain managers to make informed forecasting decisions based on average demand, demand variability, and criticality of parts to the operation (Chen et al., 2019). Claure and Ovalle (2023) show that many different schemas exist to segment MRO parts. These include policies such as separating consumable parts from spare parts, or segmenting parts based on their

annual consumption value (commonly known as ABC analysis). Other methods include XYZ analysis, which clusters parts by their demand variability, and VED analysis, which separates items by criticality – vital, essential, or desirable (Pandya & Thakkar, 2016).

2.3 Network Design and Centralization

A common strategy to reduce inventory costs in any supply chain is to hold inventory in fewer locations, a practice known as warehouse consolidation. Cost savings are realized via two channels: direct savings by having less cash tied up with inventory on hand, and indirect savings due to holding cost reductions. At an industry average holding cost of 25%, every four dollars of inventory on hand that can be reduced per year saves an additional dollar of holding costs, when considering the cost of capital as well as the opportunity cost (Azzi et al., 2014). Melachrinoudis and Min (2007) propose that through economies of scale, consolidation can reduce warehousing, inventory, and transportation costs, with potential direct savings ranging from 5% to 10%. The theory for these savings comes from the square-root rule: "if n stocking points are consolidated into one, the new inventory level I_T is related to the initial inventory I_i by $I_T = I_i\sqrt{n}$" (Melachrinoudis & Min, 2007).

A major risk when consolidating inventory is a decrease in service levels. When inventory is held in fewer locations, the average order lead time inherently increases (Melachrinoudis & Min, 2007). Milewski (2020) warns that the square-root rule "calculates only a part of stocks, namely the safety stock," and its real-world applicability is limited. Milewski (2020) also emphasizes that both safety stock and cycle stock should be included in analyses, and that simulation may provide better insight into complex systems. Because MRO supply chains are highly complex, they are well suited for simulation.

Risk-pooling is another influential concept within supply chain management. Eppen (1979) laid the foundation for the modern understanding of risk-pooling: the variability of pooled demand is less than the sum of the variability of individual demands. This means the expected holding and penalty costs in a decentralized system are greater than that of a centralized system. Berman et al. (2007) counter that the "absolute benefit of risk-pooling increases with (demand) variability" but that beyond a certain coefficient of variation, "both absolute and relative benefits decrease to zero." Their results show that centralization is robust up to a moderate degree of demand variability, but that decentralization may be preferable when demand variability is extremely high.

A resilience feature of modern network design is inventory sharing. Raised by Kutanoglu and Mahajan (2009), individual facilities within a supply chain will "share their inventory as a way to increase achieved service levels." When one facility is out of stock, demand can be met via an emergency lateral transshipment (ELT), and only when ELTs fail does the network escalate to direct delivery from the central warehouse. Sharing can both reduce the total amount of inventory held within the network and improve time-based service metrics (Kutanoglu & Mahajan, 2009). The downsides to inventory sharing are an increase in decision-making complexity and higher transportation costs due to ELTs.

Managers who want to optimize a network-type structure often use mathematical programming. These programs work by taking in a set of decision variables, a set of constraints on those variables, and an objective function that changes with the values of the variables (Caplice, 2015). Solving the program outputs a solution set of optimal value for each decision variable, which provides an optimized objective value. Mixed Integer and Linear Programming (MILP) is the sub-type of mathematical programming most commonly used in warehouse-network design problems (Caplice, 2015). As presented by Claure and Ovalle (2023), MILPs are a strong tool to design the optimal network of an MRO supply chain when the objective is to minimize total costs, given variable cost structures (product, transportation, etc.), fixed service levels, and constant or variable supply and demand levels.

2.4 Total Landed Cost Analysis

The total landed cost (TLC) of a product is calculated by summing all costs associated with procuring and transporting a product from a supplier to a customer facility (Chaudhry & Lee, 2022). By comparing the TLC of multiple different product/shipping combinations that satisfy the same need within a supply chain, companies can confidently select the lowest-cost option. In the case of our sponsor, a major variable to consider is the use of its own logistics assets and services. By performing a TLC analysis across different scenarios in its MRO supply chain, the lowest TLC options can be leveraged to broadly reduce supply chain costs.

Incoterms are a set of worldwide shipping standards used to address three critical questions:

- Who pays for freight, insurance, export duties, and unloading?
- Where does the "risk of loss or damage" transfer from the seller to the buyer?
- Who handles the paperwork, such as customs clearance and export/import licenses?

Delivered Duty Paid (DDP) is a common Incoterm within supply chains when moving MRO parts from suppliers to customers (Sponsor company, 2025). With DDP, a supplier assumes full responsibility for getting goods to a customer facility. Conversely, with the use of Free Carrier (FCA) the supplier only assumes responsibility for loading goods onto a customer's carrier, and the customer then assumes full responsibility for getting it to their own facility (International Chamber of Commerce, 2020).

When a customer wants to move material from a supplier to its own facility for a lower cost than that offered by DDP, it is presumed to be in its best financial interest to investigate switching to FCA. One way to understand if this belief is a reality is to perform a TLC analysis.

2.5 Relevance to Problem

Current work shows that managing MRO inventory is much more complex than managing traditional goods. Traditional inventory forecasting and holding policies cannot effectively be used due to the unpredictable nature of demand and the sheer magnitude of SKUs. To effectively balance inventory costs with service level requirements, techniques such as centralization and segmentation must be considered for MRO inventory. Finally, strategies such as TLC analysis and proper Incoterm selection should be used to effectively manage end-to-end supply chain costs.

3 METHODOLOGY

This section outlines the approach used to evaluate and optimize the MRO inventory network. It details the data collection and preparation processes, the multi-dimensional SKU segmentation framework used to classify inventory, and the formulation of the MILP model. Finally, it defines the progressive scenario design used to evaluate the benefits of centralization.

3.1 Methodology Selection

Based on literature review and interviews with representatives from the sponsor company, we formulated an approach combining multi-dimensional inventory segmentation with a network optimization model. Because MRO inventory spans a wide spectrum of demand patterns and criticality, applying a blanket centralization policy is suboptimal. K-Means clustering was selected to segment SKUs by spend, demand variability, and criticality, enabling tailored inventory policies. Subsequently, a MILP model was chosen to simulate and optimize the supply chain network because it can mathematically guarantee the lowest TLC while satisfying discrete facility demand and location constraints. The MILP identified optimal stocking locations, transportation modes, and material flows, while accounting for constraints such as warehouse capacities and minimum service levels. By running different input scenarios through the MILP, a comprehensive set of scenarios was developed.

3.2 Data Collection and Preparation

The sponsor provided nine discrete sheets of ERP data, as well as a list of facility and supplier locations for this project. The ERP sheets contained:

- Material master data: physical and descriptive characteristics of individual SKUs
- Bills of materials (BOMs) for every unique asset in the network: a list of the individual MRO SKUs required for each asset, as well as the criticality of each SKU to each asset
- Stocking locations and current quantities on hand, by production facility and SKU
- Purchase order data from 2020-2025, by production facility and SKU

Standardizing fragmented facility master data served as a prerequisite for network modeling. The sponsor highlighted that, due to the nature of its acquisitions, there would likely be identical items with different SKU identifiers in the ERP output. By examining the five descriptive features of each part in the database, the model performed token matching to highlight items that were likely the same physical part, with different SKU identifiers.

The raw data then underwent extensive cleaning to isolate relevant variables. Audit fields and null columns were dropped, while core inputs such as unit costs, lead times, and active bin locations were retained. Because the ERP system lacked formal safety stock parameters (recording safety stock as zero for all SKUs), baseline safety stocks were synthesized using standard normal distribution formulas based on observed demand variability and lead times. A data science pipeline was created in the Python programming language to facilitate exploration and analysis. (A pipeline is a coding framework that ensures consistent, expected handling of data for analysis. This pipeline ingested all the data and established a current-state baseline for initial network design analysis). Finally, geographic routing distances between suppliers and facilities were calculated using the SCGraph road-network dataset, providing the foundational transport parameters for the MILP model.

3.3 Application of Methodology

3.3.1 Multi-Dimensional SKU Segmentation

Following data standardization, the model categorized SKUs using a multi-dimensional segmentation framework adapted from Jackson (2022). This approach was selected because relying on a single dimension, such as historical spend, is insufficient for MRO environments where the absence of critical, low-spend parts can halt production. Therefore, the model segmented the data across three dimensions:

- ABC Analysis: A items are the 15% of SKUs that make up 80% of total spend, B items are the 20% of SKUs that make 15% of total spend, and C items are the remaining 65% of SKUs that make up 5% of total spend.
- XYZ Analysis: X items have low demand volatility; Y items have moderate demand variability; Z items have high demand volatility.
- VED Analysis: The sponsor uses SKU criticality codes, C1-C9. C1-C3 are mapped to desirable for operations, C4-C7 are mapped to essential for operations, and C8-C9 are mapped to vital for operations.

These dimensions were fed into a K-Means clustering algorithm, resulting in distinct inventory policy clusters: *P1 (Centralize Now)*, *P2 (Centralize with Buffer)*, and *P3/P4 (Keep Local)*. This ensured that only high-spend, high-criticality items with centralization potential (P1/P2) were routed into the MILP network design, while low-value, sporadic items (P3/P4) remained locally managed. The top 12 clusters are shown visually with a heatmap in Figure 1.

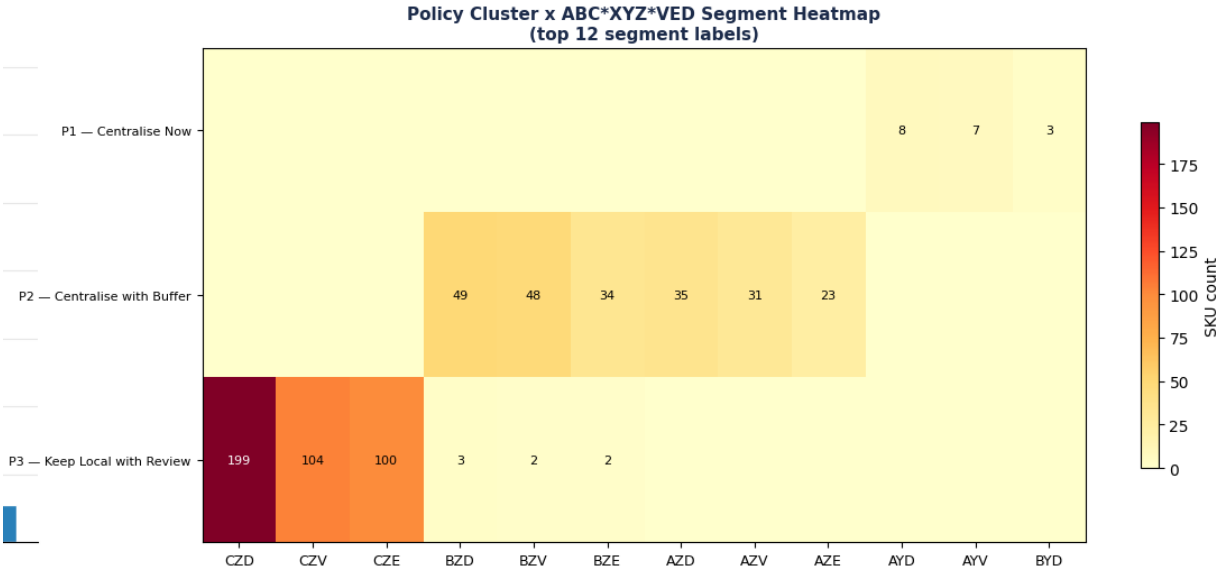


Figure 1. Heatmap of Top 12 Inventory Policy Clusters

3.3.2 MILP Model Design

A multi-product MILP model was formulated to identify the optimal hub locations and item flows. The objective was to minimize network TLC.

The objective function minimizes the sum of inbound transportation, outbound transportation, carrying costs, and fixed hub costs:

$$Minimize\ Z = \sum_{s,h,i} c_{shi} X_{shi} + \sum_{h,f,i} c_{hfi} Y_{hfi} + \sum_{f,i} h_{fi} I_{fi} + \sum_{h} F_h W_h$$

Where:

- c_{shi} is the unit transport cost from supplier s to hub h for item i.
- X_{shi} is the flow volume from supplier to hub.
- c_{hfi} is the transport cost from hub h to facility f.
- Y_{hfi} is the outbound flow volume.
- h_{fi} is the holding cost per unit (calculated at a 25% annual carrying rate).
- I_{fi} is the expected inventory level (cycle stock plus safety stock, adjusted by the square root law of pooling for centralized SKUs).
- F_h is the fixed operational or designation cost of opening hub h.
- W_h is a binary decision variable representing whether hub h is open.

The model was subjected to flow conservation constraints, single-sourcing constraints (no cross-border flows were permitted between the United States and Canada), and demand fulfillment constraints.

3.3.3 Network Routing and Internal Logistics Simulation

To accurately simulate freight costs and lead times in the subsequent MILP model, the model generated a highly accurate geographic distance matrix between the given suppliers and facilities.

Traditional Haversine calculations have often overestimated corridor distances by large margins, even when adjusted for circuity. To overcome this, the model utilizes the SCGraph package (Makowski et al., 2025). This tool uses a shortest-path routing algorithm on global OpenStreetMap highway networks to calculate true road network distances. The resulting supplier and facility distance matrices map a geographic spread of approximately 700 kilometers across the supply chain. This precise routing data serves as the primary input for evaluating the feasibility of internal logistics utilization in the optimization phase.

3.3.4 Scenario Design

To isolate the drivers of cost savings, the MILP was evaluated across progressive scenarios:

- Scenario A: Decentralized Baseline (current state of 15 independent nodes).
- Scenario B & C: Fully Centralized and Hybrid Centralized (testing pure centroid-based routing).
- Scenario D: Optimal Leased Hub (MILP selects an optimal commercial location assuming a $75,000 annual lease).
- Scenario E & F: Existing Facility Hubs (forcing the hub to be an existing property, avoiding heavy capital expenditure).
- Scenario G: Optimal Number and Location of Existing Hubs (MILP freely selects any current facility as a hub based on a scaled hub designation cost).

4 RESULTS AND DISCUSSION

This section presents the outcomes of the data segmentation and the network optimization scenarios. It details the quantitative benefits of centralization, provides a rationale for the recommended supply chain architecture, and discusses the limitations of the analysis.

4.1 Current State and Segmentation Results

The current state baseline scenario revealed a highly fragmented procurement and stocking strategy. Across all active facilities, the network processed 7,982 purchase order lines over five years, totaling $1,482,906 in five-year MRO spend. Six hundred and seventy-four unique SKUs were held in active inventory, totaling 6,006 units on hand in the network. A simplified network design model of the current state is shown in Figure 2.

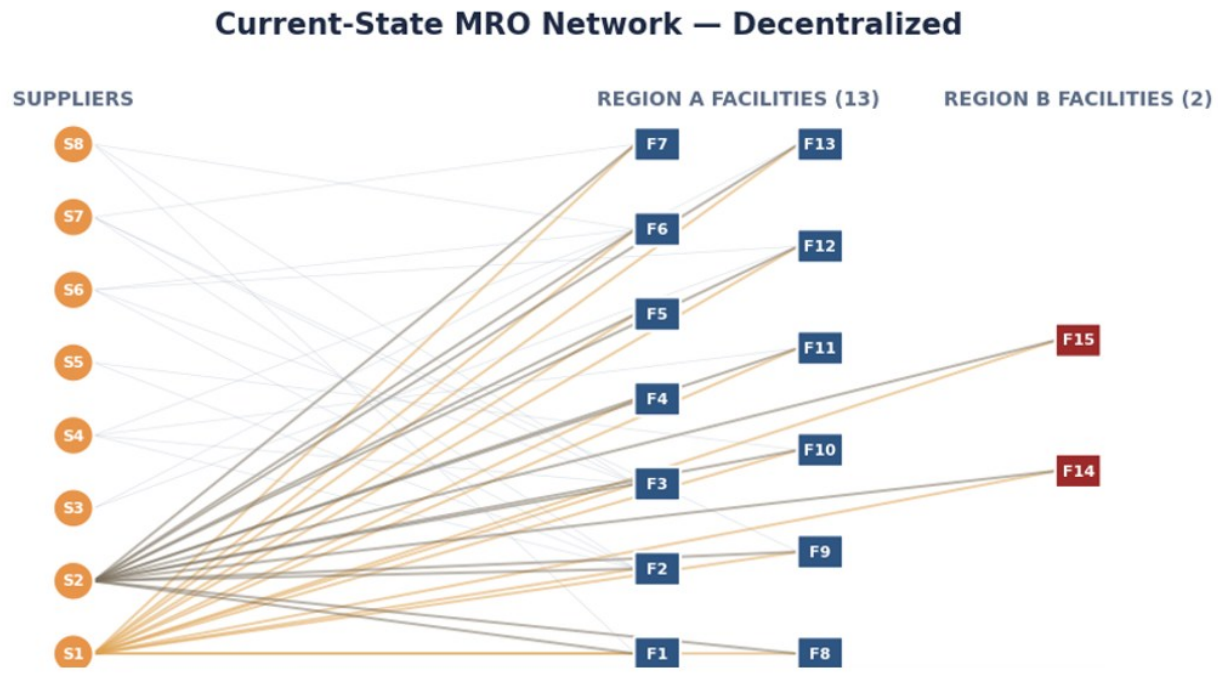


Figure 2. Current State Baseline

Aggregating MRO consumption by facility revealed significant variance in MRO consumption across the network. A single facility represented the highest individual spend at $329,515. Analysis of historical purchase orders revealed that local facility managers ordered and stored inventory independently of one another, even when procuring identical categories of materials from shared regional vendors. Because individual facilities maintain disjointed warehousing data, the business lacks the visibility to transfer stock between facilities to mitigate downtime events. This lack of visibility leads to inflated safety stock levels at each site. The total five-year spend of nearly $1.5 million across just 674 active SKUs highlights a dense concentration of capital that is tied up in maintenance redundancies.

The algorithm successfully identified duplicate and highly similar parts, which were stocked under many different internal SKU identifiers. Figure 3 shows 59 total groupings of items that contain two or more SKU identifiers for the same component.

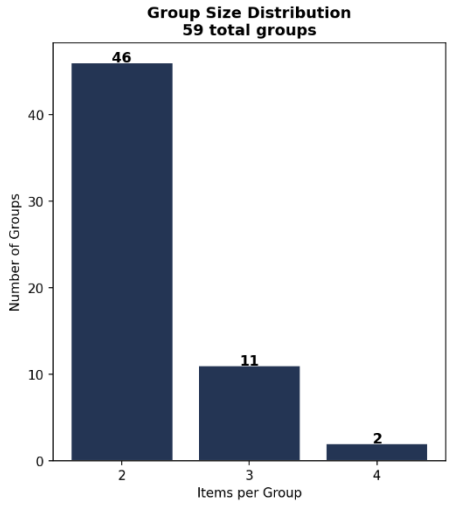


Figure 3. Duplicate SKU Groupings

The multi-dimensional K-Means segmentation successfully triaged the inventory. Exactly 241 unique SKUs (representing 94% of the $296,581 annual MRO spend) clustered into P1 and P2 groups. These vital, high-spend items became the focus of the centralization model. The remaining 410 SKUs (representing only 6% of spend) were assigned to P3 and P4 clusters and were excluded from centralization, as routing low-value, sporadic items through a hub would incur transit penalties without realizing meaningful pooling benefits.

The application of string matching and the multidimensional segmentation framework fundamentally changes how the sponsor should view MRO inventory. The addition of volatility and criticality metrics to traditional spend-only inventory segmentation ensures that the proposed centralized network will not remove operationally critical components from being held at each mill. For example, a vital but rarely used bearing may have a low spend and high volatility. Traditional metrics make it inefficient to hold this item. However, the proposed segmentation ensures that the system flags this item as having a high downtime risk, requiring the facility to maintain a decentralized safety stock. Conversely, fast-moving, noncritical consumables identified by the K-Means algorithm are ideal candidates for a centralized hub.

4.2 Network Optimization Findings

A comparison of savings realized by Scenarios A through G is shown in Figure 4. Ultimately the MILP model of Scenario G demonstrated that moving from a fully decentralized network to a hybrid centralized network yields especially significant financial and operational benefits.

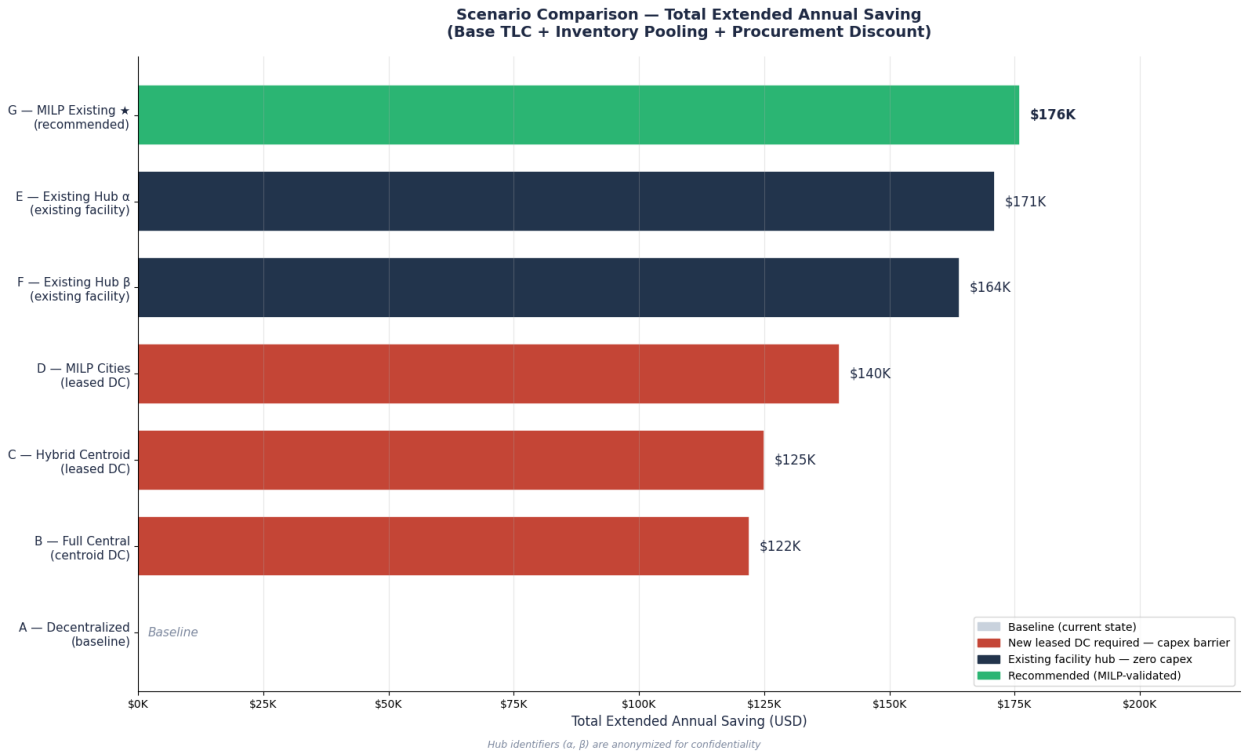


Figure 4. Scenario Comparison — Total Extended Annual Saving Across All Seven Network Designs

Scenario G was utilized to find the optimal hub architecture utilizing existing facilities. Through sensitivity analysis, it was discovered that at an operational designation cost of $10,000 per hub, the model mathematically converges on two hubs one for the Canadian network and one for the U.S. network.

The model selected the Canadian hub primarily because it is the center of gravity for MRO consumption; it has the highest historical order volume as well as the highest five-year spend ($329,515). Selecting this site minimizes outbound transportation costs to the remaining Canadian facilities. Similarly, the U.S. site was selected to minimize total geographic routing distances. A simplified network design model of the Scenario G state is shown in Figure 5.

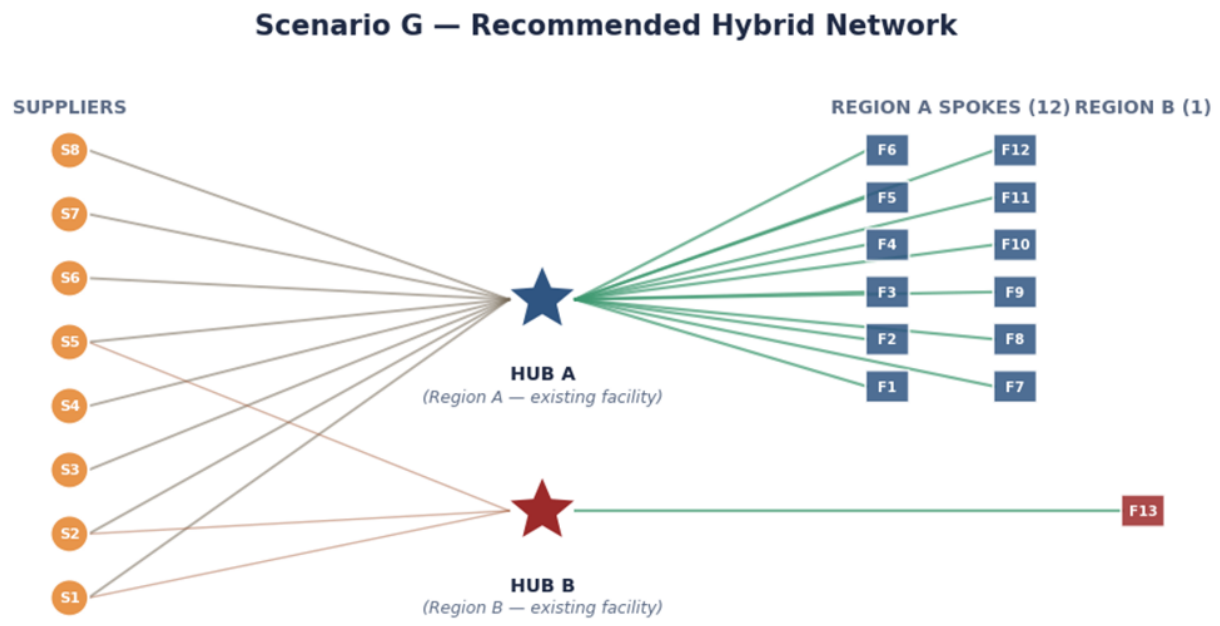


Figure 5. Scenario G

4.3 Inventory and Cost Implications

The annual cost of the current state baseline (Scenario A) is $211,517, comprising $13,979 in carrying costs, $17,538 in transport costs, and an estimated $180,000 in yearly storeroom operational overhead. The overhead assumption reflects the marginal cost of processing the given subset of MRO SKUs through the 15 independent storerooms, estimated at $12,000 per location per year to cover operator time, WMS licensing, cycle counting, and procurement administration. By transitioning P1 and P2 items to the hubs, the total landed cost dropped to $140,851, which is a structural base savings of $70,666 annually. Beyond the base TLC saving, centralization unlocks two additional recurring saving layers: a $90,473 annual carrying cost reduction from inventory pooling (detailed in Section 4.6), and a conservative $14,829 annual procurement saving from consolidating 15 independent purchase orders into a single program (estimated at 5% of the $296,581 annual spend). Combined, the total extended annual savings are approximately $175,968. In addition, the consolidation of 6,006 distributed on-hand units to a pooled requirement of 1,666 units at the hubs represents a one-time capital release of $361,894. This is calculated as 4,340 freed units at an average unit cost of $83.38. A detailed breakdown for Scenario G is shown in Figure 6.

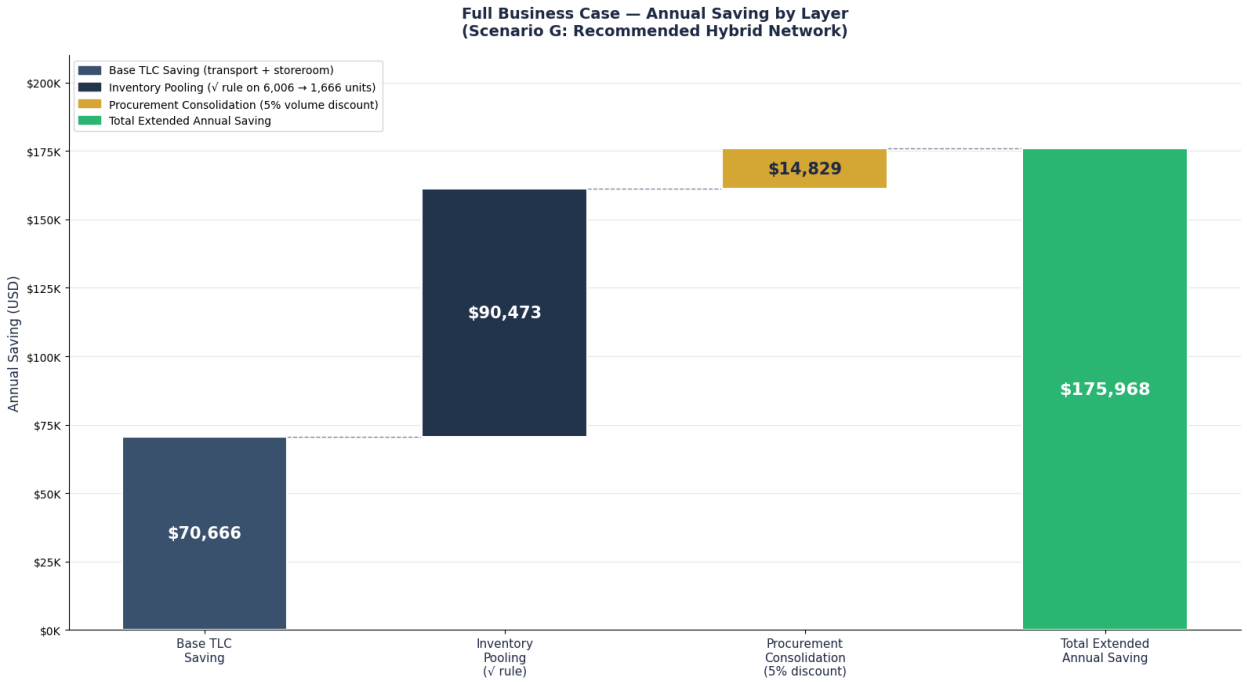


Figure 6. Full Business Case — Annual Saving by Layer (Scenario G)

A primary driver of these savings was the reduction in safety stock achieved through risk pooling. By holding inventory centrally, the network benefits from the square root law of inventory management. The MILP outputs showed that for the centralized SKUs, required safety stock units dropped by 33% in Canada (from 552 to 368 units) and 12% in the United States (from 70 to 62 units).

4.4 Sensitivity Analysis: Hub Robustness and Designation Costs

To ensure the recommended network is resilient to fluctuations in operational costs, a sensitivity analysis was performed on Scenario G. The model was tested against a range of hub designation costs that represent the incremental overhead of managing a centralized facility, from $0 to $15,000 per year per hub.

At a $0 designation cost, the MILP naturally gravitates toward a transport-minimum solution, opening hubs at all 15 facilities to minimize outbound distances. However, as the designation cost increases to $10,000 per year, the network architecture converges and stabilizes at exactly two hubs. This convergence suggests that the selection of the hubs is not merely a result of cost-minimization but is driven by their strategic positioning relative to high-volume SKU consumption. Even as coordination costs increase by 50% (to $15,000 per hub), the two-hub topology remains optimal, indicating a high degree of structural robustness in the proposed design. This analysis is shown graphically in Figure 7.

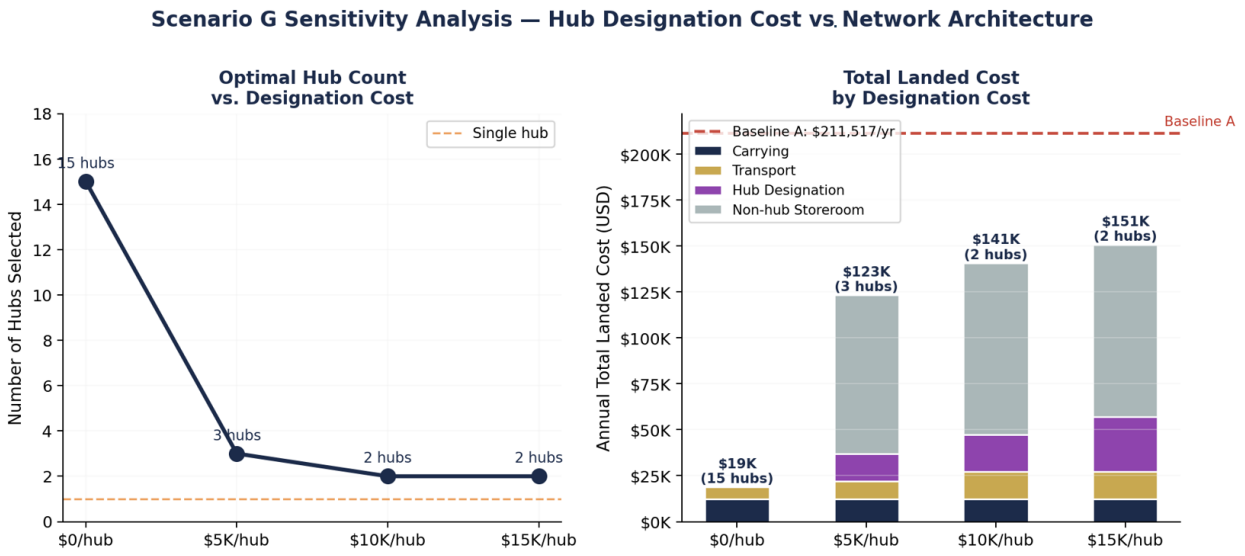


Figure 7. Scenario G Sensitivity Analysis — Hub Designation Cost vs. Network Architecture and Total Landed Cost

4.5 Predictive Modeling Performance and Demand Characterization

The use of machine learning, specifically XGBoost, to forecast MRO demand addresses a significant challenge in inventory management: intermittency. Analysis of the sponsor's five-year purchase history revealed that approximately 85% of active SKUs are classified as "Z-class," characterized by a Coefficient of Variation (CV) greater than or equal to 1.0, with the remaining 15% classified as X or Y class based on lower demand volatility.

While traditional time-series models often fail to capture the "lumpy" demand of items like conveyor belts or spherical roller bearings, the gradient-boosted tree approach allows for the simultaneous integration of temporal features and SKU-specific attributes. By utilizing a log-transformed spend variable for SKU clustering, the segmentation remains stable even when outliers are present. This stability is critical, because it ensures that the MILP model is fed with reliable demand clusters that do not shift erratically due to single-period spikes in procurement.

4.6 Inventory Policy Dynamics: Risk Pooling and the Square Root Law

The quantitative benefits of centralization are most visible in the reduction of safety stock levels. By transitioning the P1 and P2 policy clusters from 15 independent facilities to two centralized hubs, the network realizes the mathematical benefits of the square root law of inventory pooling.

For the centralized SKU subset, the total safety stock units required to maintain the desired service levels (95%–99%) decreased from 622 units (552 in Canada plus 70 in the United States) in the decentralized state to 430 units (368 in Canada plus 62 in the United States after centralization). This represents a 31% reduction for the centralized items. Including the 162 units of P3/P4 safety stock that remain locally held at each facility, the total network-wide safety stock moves from 784 units to approximately 592 units, a 24.5% reduction overall. This reduction in physical units held translates to a significant release of working capital. This efficiency is particularly vital for MRO items where the cost of a stockout far exceeds the incremental transport cost of centralized delivery.

4.7 Limitations

While the results strongly advocate a hybrid centralized model, several assumptions influence the precision of the output:

1. Transport Rate Assumptions: The model utilizes a baseline transport rate of $0.012 per unit-km derived from regional Less-than-truckload, or LTL, benchmarks. If the sponsor's internal fleet operates at a vastly different cost scale, the transport penalty for hub routing may shift.
2. Carrying Cost Estimations: Because the ERP system recorded zero carrying costs and zero safety stock, an industry-standard 25% annual carrying rate was applied. Variations in actual capital opportunity costs could alter the TLC calculations.
3. Routing Constraints: Distances were calculated using the SCGraph road network. The model strictly prohibits cross-border flows between the United States and Canada to simplify customs and tariffs, though real-world operational exceptions may occasionally occur.
4. Incoterms: The network model evaluates structural logistics but treats the FCA versus DDP procurement switch as a parallel opportunity. The exact realization of inbound freight savings depends on active supplier renegotiation. Coupled with the fact that available ERP data doesn't capture freight costs in a way that makes comparison possible, a detailed analysis was not performed.

5 RECOMMENDATIONS

The primary objective of this project was to evaluate whether a more centralized MRO inventory structure could maintain high service reliability while reducing inventory and transportation costs for a large, vertically integrated conglomerate. By pairing a multi-dimensional SKU segmentation framework with a MILP network optimization model, the study identified significant inefficiencies in the current decentralized state. Key findings demonstrate that a hybrid centralized architecture anchored by two regional hubs can reduce TLCs by 33% through structural savings and the mathematical benefits of risk pooling.

5.1 Management Recommendations

The following tactical and strategic recommendations offer a pathway for the sponsor to act on these findings to enhance MRO efficiency, cost transparency, and operational scalability.

5.1.1 Immediate Tactical Actions: Master Data and Segmentation

To unlock economies of scale in procurement and streamline master data cleanliness, the sponsor should immediately industrialize SKU rationalization efforts into its ERP. This action addresses the 59 identified groupings of duplicate and highly similar parts that are currently stocked under different internal SKU identifiers. By conducting routine automated audits of vendor part numbers, the organization can prevent future master data proliferation and reduce the capital that is tied up in redundant inventory.

Furthermore, facility managers should transition from intuition-based safety stock parameters to policies governed by the ABC, XYZ, and VED clusters established in this study. Implementing automated minimum reorder points for "vital" criticality items and strict consolidation for "desirable" items will ensure that service levels are maintained while minimizing unnecessary stock at local branches.

5.1.2 Strategic Network Redesign

The sponsor should adopt the hybrid centralized network design (Scenario G). By centralizing the 241 P1 and P2 SKUs which represent 94% of the total annual MRO spend, the organization can leverage the square root law of risk pooling to reduce required safety stock units by 33% in Canada and 12% in the United States.

This strategic redesign provides approximately $175,968 in total extended annual savings: $70,666 in base TLC savings from structural network improvements, $90,473 in inventory carrying cost savings from risk pooling, and $14,829 from procurement consolidation at a conservative 5% volume discount, reducing the network's TLC from a baseline of $211,517 to $140,851. In addition, the transition releases $361,894 in one-time working capital that is currently locked in redundant decentralized safety stock. These savings directly support corporate strategic goals of profit growth and operational efficiency.

5.1.3 Leveraging Vertical Integration for Internal Logistics

To further reduce total landed costs, the sponsor should leverage its existing transportation network to implement scheduled, regional short-haul multi-destination routes originating from the new centralized MRO hubs. By shifting procurement terms with major regional suppliers from DDP to FCA, the organization can leverage its internal fleet capacity to pick up bulk MRO orders.

This shift eliminates vendor freight markups and allows the company to capture the hidden margin within freight costs. Internal dispatchers should utilize the real-world road network matrices generated via the SCGraph package to optimize these regional routes across the Canadian and U.S. facilities.

5.2 Future Work

While this project establishes a robust mathematical foundation for MRO network design, the organization should consider the following next steps to expand on these accomplishments:

- Predictive Maintenance Integration: As data maturity increases, the organization could integrate predictive maintenance models utilizing Internet of Things (IoT) sensor data from mill machinery into the forecasting pipeline to shift from reactive to predictive procurement.
- Stochastic Modeling: Future researchers could expand the MILP model into a multiperiod, stochastic formulation to account for severe seasonal weather variations that frequently affect road networks and operations in the Canadian and U.S. corridors.
- Strategic Scalability: The value of these findings extends beyond the project sponsor to the broader industrial MRO industry, demonstrating that mathematical optimization and multi-dimensional segmentation can provide a resilient template for managing complex, intermittent demand in vertically integrated conglomerates.

REFERENCES

Azzi, A., Battini, D., Faccio, M., Persona, A., & Sgarbossa, F. (2014). Inventory holding costs measurement: a multi-case study. *The International Journal of Logistics Management*, 25(1), 109-132.

Bacchetti, A., & Saccani, N. (2012). Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice. *Omega*, 40(6), 722-737.

Berman, O., Krass, D., & Mahdi Tajbakhsh, M. (2011). On the benefits of risk pooling in inventory management. *Production and Operations Management*, 20(1), 57-71.

Caplice, C. (2015). SC2x-Supply Chain & Logistics Design – *Intro to Network Design*. Center for Transportation and Logistics.

Chaudhry, M. S., & Lee, K. Z. (2022). *Total Landed Cost Model*. Massachusetts Institute of Technology.

Chen, J., Gusikhin, O., Finkenstaedt, W., & Liu, Y. N. (2019). Maintenance, repair, and operations parts inventory management in the era of industry 4.0. *IFAC-PapersOnLine*, 52(13), 171-176.

Claure, M. A. T., & Ovalle, T. R. C. (2023). *Network Design in MRO Inventory for Oil & Gas Company*. Massachusetts Institute of Technology.

Demke, T. M., Kämpfer, T., Lucht, T., Wachsmann, J., & Nyhuis, P. (2023, September). Forecast-Based Dimensioning of Spare Parts Inventory Levels in the MRO Industry. In *IFIP International Conference on Advances in Production Management Systems* (pp. 92-106).

Demke, T. M., Osterkamp, L., Stephan, M., Wachsmann, J., Nyhuis, P., & Schmidt, M. (2025). Investigation of different forecasting methods and models using a Forecast Support System to forecast spare parts demand in the MRO industry. *Procedia CIRP*, 135, 213-220.

Eppen, G. D. (1979). Note—effects of centralization on expected costs in a multi-location newsboy problem. *Management Science*, 25(5), 498-501.

International Chamber of Commerce. (2020). Incoterms® 2020: Rules for the use of domestic and international trade terms. *ICC Services Publications Department*.

Jackson, I. (2022). AutoML Approach to Stock Keeping Units Segmentation. *Journal of Theoretical and Applied Electronic Commerce Research*, 17(4), 1512-1528.

Kutanoglu, E., & Mahajan, M. (2009). An inventory sharing and allocation method for a multi-location service parts logistics network with time-based service levels. *European Journal of Operational Research*, 194(3), 728-742.

Melachrinoudis, E., & Min, H. (2007). Redesigning a warehouse network. *European Journal of Operational Research*, 176(1), 210-229.

Makowski, C., Saragih, A., Guter, W., Russell, T., Heinold, A., & Lekkakos, S. (2025). SCGraph: A dependency-free Python package for road, rail, and maritime shortest path routing generation and distance estimation. MIT Center for Transportation & Logistics Research Paper Series.

Milewski, D. (2020). Total costs of centralized and decentralized inventory strategies—including external costs. *Sustainability*, 12(22), 9346.

Pandya, B., & Thakkar, H. (2016). A review on inventory management control techniques: ABC-XYZ analysis. *REST Journal on Emerging Trends in Modelling and Manufacturing*, 2(3), 2016.

Willemain, T. R., Smart, C. N., Shockor, J. H., & DeSautels, P. A. (1994). Forecasting intermittent demand in manufacturing: a comparative evaluation of Croston's method. *International Journal of Forecasting*, 10(4), 529-538.