Network Design in Maintenance Inventories for Electric Utilities

by

Jason Andrew Maen
Bachelor of Science Supply Chain Management, Bachelor of Science Accounting
Iowa State University - Ivy College of Business, 2017

And

Nicholas Shiverick Winters
Bachelor of Arts in Political Science and Economics
Dalhousie University, 2009

SUBMITTED TO THE PROGRAM IN SUPPLY CHAIN MANAGEMENT
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF APPLIED SCIENCE IN SUPPLY CHAIN MANAGEMENT
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

Signature of Author: ____

Department of Supply Chain Management
May 6, 2022

Signature of Author: ____

Department of Supply Chain Management
May 6, 2022

Certified by: ____

Tim Russell, Research Engineer, Humanitarian Supply Chain Lab
Capstone Advisor

Accepted by: ____

Prof. Yossi Sheffi
Director, Center for Transportation and Logistics
Elisha Gray II Professor of Engineering Systems
Professor, Civil and Environmental Engineering

Network Design in Maintenance Inventories for Electric Utilities

by

Jason Andrew Maen

and

Nicholas Shiverick Winters

Submitted to the Program in Supply Chain Management
on May 6, 2022, in Partial Fulfillment of the
Requirements for the Degree of Master of Applied Science in Supply Chain Management

Abstract

Within the electric distribution space maintaining a balanced inventory of spare parts forms a critical component of resiliency and restoration in the event of an outage. We find that leveraging network design provides an opportunity for utility companies to improve the effectiveness of the inventory they hold, enabling better service at a lower cost. For one regional utility company in the United States, an inventory reduction of 35% was found by adopting a hub and spoke model over a previous decentralized model. We believe that this observation can be extended to other companies with distributed assets, highly variable demand for inventory, long lead times, and a high cost of downtime.

Capstone Advisor: Tim Russell

Title: Research Engineer, Humanitarian Supply Chain Lab, Center for Transportation and Logistics

Acknowledgements

Joint Acknowledgments:

We would like to thank our advisor Tim Russell for the support he provided to us in helping us frame our thoughts and tackle this problem more fully. We would like to thank our writing coach Pamela Siska for helping us work through our writing and improve it.

We also thank our sponsor and their team members for helping us gather the data and working with us to understand the underlying structure and format; as well as discussing with us the challenges and opportunities on which our analysis relied.

We would also like to thank the supporting companies who supplied us with the software that we used to complete our analysis including Coupa for supplying us licenses to Supply Chain Guru to manage and frame our network models and to Alteryx for supplying us with Designer licenses we used to examine and clean our data. We would also like to thank our classmates for their insights and support throughout this journey.

Jason Maen Acknowledgments

I thank my partner for his instrumental work throughout this project and for the insights he helped me to unlock. Additionally, my close network of supporters back home who helped along the way with notes of encouragement and care packages of refreshments for late nights. Finally, to my wife, Bethany, my deepest gratitude for your support, and appreciation for setting me on this path and believing in me when I did not see the way through. Your selflessness and sacrifice enabled this journey and your love and support carried me through, thank you.

Nicholas Winters Acknowledgements:

I would like to thank my wife Mallory for all of the support through this journey, and to thank my capstone partner for his assistance in completing this project.

Table of Contents

Abstract 2

Acknowledgements 3

Table of Contents 4

List of Figures 5

List of Tables 6

1.0 Introduction 7

1.1 Sponsor Background: 8

2.0 Literature Review 9

2.1 Impact of MRO Inventories on Companies 10

2.2 Strategies for Managing MRO Inventories 10

2.3 MRO Inventory Management in the Electric Utilities Market 13

2.4 Key Lessons 15

3.0 Methodology 16

3.1 Data Gathering 17

3.2 Inventory Optimization Model Structure 19

3.3 Demand Analysis 22

3.4 Validity Check 25

3.5 Detailed Model Construction 28

4.0 Results and Analysis 35

4.1 Baseline Comparison 37

4.2 Impacts not modeled 39

5.0 Discussion 43

5.1 Implementation Strategy 45

5.2 Storm Stocks 46

6.0 Conclusion 47

References 48

List of Figures

Figure 1: Histogram of spend by project 18
Figure 2: Vendor distribution 20
Figure 3 Demand by lead time 21
Figure 4: Demand classification flow chart (Coupa, 2021) 23
Figure 5: Demand summary across categories 24
Figure 6: Decentralized Supply Construction 31
Figure 7 Hub and spoke framework 32
Figure 8 Lead Time implementation in hub and spoke and decentralized models 34
Figure 9: Impact of network design on required safety stock 36
Figure 10 Mathematical relationship between cycle service level and full order rate assuming independent items 38

List of Tables

Table 1: Demand classifications in centralized and decentralized models 25
Table 2: Counterfactual, a simple reduction 26
Table 3: Demand over lead time for storerooms under decentralized and centralized inventory 26
Table 4: Changes in Storeroom safety stocks from faster replenishment associated with a hub replenishment 28
Table 5: Summary of impacts not modelled 40
Table 6: Expected Savings at 99% cycle service level 43

1.0 Introduction

Worldwide, $2.6 trillion dollars USD is spent on the electric utilities market annually (IEA, 2000-2020) with some markets estimating that as much as one-third of that expenditure for residential consumers occurs within electric distribution (Alberta Energy System Operator, 2020) (US Energy Information Administration, 2020). While most supply chains have inventories, electricity requires instantaneous alignment between the electricity generated and the electricity consumed, and this delicate balance must be maintained through continuous uptime of the grid. When electrical grids fail to supply power, this can have serious and immediate costs to consumers. During 2021's Winter Storm Uri, which struck the central and southern US, deaths were linked to the loss of power limiting the ability to heat homes or to provide necessary medical care (Hauser & Sandoval, 2021). While some customer sites carry sufficient power to safely power down their operations in the event of an outage, very few locations carry sufficient electricity storage to operate during an outage. It is unsurprising that reliability forms the second most common complaint among dissatisfied electric customers in the United States behind cost (S&C, 2018).

A major challenge for power companies is to improve their ability to maintain service to their customers while operating within their budgets. To maintain service, power companies in the US must not only maintain generation facilities but also 5.7 million miles of transmission and distribution infrastructure (US Energy Information Administration, 2021; Weeks, 2010).

The project sponsor, a regional power company integrated across generation, transmission, and distribution, maintains approximately $136 million dollars of spare parts inventory of which $50 million dollars in inventory to support its electric distribution operations. They have engaged us to identify a simpler, more efficient, and more reliable central stores system to support their power distribution business.

We test whether power distribution companies can improve service at a lower cost, by utilizing a hub and spoke model over a decentralized inventory model for their power distribution inventories. Our null hypothesis is that a hub and spoke model offers no improvement to cost or service over a decentralized model.

1.1 Sponsor Background:

The sponsor currently manages electric distribution in three[1] states. To support their service area, they currently operate a decentralized network of 24 major storerooms and a total network of 72 storerooms. These storerooms are supported by a central purchasing team who source the parts each storeroom requires, either externally or through lateral shipments between storerooms.

The sponsor has identified limitations in this approach and in response this capstone seeks to determine whether a hub and spoke model, where a central warehouse orders, receives, batches, and distributes materials to the storerooms, may drive lower costs, simplify execution, as well as improve service for the crews tasked with repairing and constructing its distribution infrastructure.

We will conduct a detailed examination of the impacts of a hub and spoke model, including severe weather events, which are viewed to be the "true test of any solution" (Logistics & Support, 2021)[2]. Our research will supply the sponsoring entity with a specific answer to their business context, but also supply a broader understanding for other companies who observe similarities within their own organization.

Many companies have looked to ensure that their operations are reliable and enduring. Often these initiatives fall into projects to ensure that their Maintenance Repair and Operations (MRO)

---

[1] Number anonymized
[2] These are the anonymized titles of the interviewees. Names are redacted to protect confidentiality of the sponsor. Throughout this paper details which might be used to identify the sponsor have been anonymized

inventory is effectively planned and utilized. Our literature review reveals that despite extensive writing on this topic there is minimal writing on this area within the electric utilities space and none examining a hub and spoke model. By creating a model of this network, we will fill a gap in the current literature and create a framework that could be of direct use to a market which represents approximately 0.6 percent of the US's GDP[3], or approximately three times the size of Amazon.

2.0 Literature Review

Many organizations segment their purchasing strategies between their direct materials, the items that they consume in the performance of their function, and their Maintenance, Repair, and Operations (MRO) items. The MRO category often consists of a much larger collection of less frequently used items that are not consumed in the primary function but are necessary to the ongoing functioning of the organization. In electric distribution the materials for repair and the expansion of the network are the same materials. In this context the literature on MRO materials provides an insight to how other companies and industries have solved the problem of setting inventory levels for areas with a large number of parts, where downtime is a primary concern.

This chapter reviews the literature on MRO inventories. First, we will examine the context of MRO inventories within companies and why this has formed a separate section of supply chain literature. Second, we will review the most common strategies that companies have taken to improve the performance of their MRO supply chains. Third, we will engage in a specific examination of the

[3] US retail expenditures on electricity in 2019 were $399 billion dollars (US Energy Information Administration, 2019) distribution represents 31% of the average costs of electricity (US Energy Information Administration, 2020) equating to $124 billion. US GDP in Q4 2019 was $21,694 billion (U.S. Bureau of Economic Analysis). Note this does not include indirect impacts, only direct expenditures.

literature within the context of the power industry broadly. Finally, we will summarize the key lessons as they apply to this study drawing from both the power industry and the broader literature.

2.1 Impact of MRO Inventories on Companies

Many organizations carry millions to billions of dollars in spare parts (Basten & van Houtum, 2014). Long term trends in industry, particularly the inclusion of just in time strategies for direct materials have amplified the importance of these spare parts inventories (Gilbert & Finch, 1985). MRO inventories are also growing due to the increasing technological complexity of equipment (Ghodrati & Kumar, 2005). The trends of increased importance, increased volume, and increased cost have led many researchers to investigate methods to improve MRO inventories.

A Google Scholar search yielded over 11,000 papers on the subject of MRO inventory. Despite the available research, MRO inventory management continues to pose a challenge for companies. Bechtel and Patterson (1997) state: “Despite the increased attention procurement managers are giving MRO purchasing, success in managing MRO items has been limited.”

An extreme example of the challenges Gilbert and Finch (1985) identified with just in time inventories increasing the importance of MRO, is the sphere of electric distribution. With electricity, storage is limited, and alignment between generation and distribution must be near instantaneous. Any downtime, therefore, causes an immediate loss of revenue for the utility and immediate consequences for their customers.

2.2 Strategies for Managing MRO Inventories

In both the literature and the policies of the sponsor, we see attempts to improve inventory management by implementing approaches based on standard supply chain strategies. We have identified three common approaches for handling the MRO inventory challenge. As a common first tactic, companies engage in classifying and segregating inventory to create rules and focus management attention. The second common approach is to improve the forecasting of demand. A third method is to examine ways of pooling demand, by looking at the MRO supply chain as a network. This network can be within a firm or between firms. In the context of electric distribution, we are focusing on internal network design.

Creating a segregation and categorization of inventory is important as a first approach for companies, because it identifies critical items that require management attention. Bechtel & Patterson (1997) do this by recommending the formation of commodity groups that can lead to strategic partnerships and improve the overall performance of the MRO inventory, whereas Schroder (2004) breaks these down by factors of criticality, lead time, and historical usage. We noted in our interviews with the sponsor that they had already adopted a variety of these approaches, from commodity groupings to conventional ABC analysis, to evaluating their system for caches of "reserved" inventory that are held for "what-if" scenarios. Further, while the sponsor had initially seen success in reducing the inventory levels, we noted that the identification and procurement of critical items led to the return of previous inventory valuation levels (Logistics & Support, 2021)[4].

Second, improving demand forecasting is a standard approach to handle inventory management and MRO is no exception. Mann (1966) provided an initial summary of this strategy, proposing that companies can provide a general estimate of a lifetime of a unit and predict its failures. This predictive modeling limits the amount of unplanned downtime for an asset and is an ongoing area of research in the distribution utility segment. Advancements to this model have involved the inclusion of proportional hazards models to better estimate failures. This includes the estimation of failures at the start due to design installation or defect, as well as failures towards the end of service life or based on operating

[4] Names are redacted to maintain confidentiality of the sponsor, functional roles are referenced instead

conditions. Advances in technology, including increased computational power, decreasing costs of sensors, and improvements in predictive analytics, have resulted in renewed interest in predictive maintenance, as part of the larger "Industry 4.0" trend. Chen et al. (2019) point to how these improved analytics and sensors can improve the forecasting of demand. Forecasting provides a potential future avenue for improvement for the sponsor, by improving datasets around asset age and asset condition to predict failures. The ability to predict failures allows better allocation and forecasting of inventory and in turn decreases or potentially prevents downtime while allowing inventory costs to be lowered.

The third common strategy involves examining the supply chain as a network to optimize the flow of goods between facilities. This examination can be done either within a company as a flow of goods between warehouses, or between companies, as they attempt to optimize logistics flows and pool risk. The seminal examination of this problem was Sherbrooke's (1966) METRIC model. The METRIC model was initially an examination of US Air Force repair policies when there might be an airbase with limited repair capabilities as well as a central depot which could also repair the aircraft. The model attempted to minimize backorders for parts and, in turn, manage both parts and aircraft availability. This model was significantly advanced in the following years. Muckstadt (1973) introduced the MOD-METRIC extension in 1973 to create a systems-based view of part prioritization. Further developments to the METRIC model, including VARI-METRIC, were shown to result in models which were highly accurate. The accuracy of these models was within 1% of optimality while remaining computationally efficient (Sherbrooke, 1986). These developments of the model have resulted in a mature approach that can be leveraged in new spaces as we will do with the sponsor's network analysis.

This final strategy of examining a network closely fits with the sponsor's question regarding their inventory structure. Specifically, how to organize their current warehousing strategies and whether a hub and spoke model, similar to the initial item modeled in METRIC, would provide an improvement over a decentralized approach. Further, it will allow us to examine existing literature which has already been adopted across an array of fields, to a new industry.

2.3 MRO Inventory Management in the Electric Utilities Market

Limited research is available on the supply chains of electric distribution. This stands in contrast to the relatively extensive MRO research, and to the size of the power market. Out of the 11,000 previously stated Google Scholar articles, only four major papers were identified in this field. Each of these has focused on three traditional methods of inventory management: segmentation, demand forecasting, and network analysis. We will discuss each in turn as well as the areas of impact.

Bailey and Helms (2007) described an approach to reduce the MRO inventories within power generation facilities of the Tennessee Valley Authority. While Bailey and Helm's focused on an allied field, they outlined the broader issues with managing spare parts in the electric generation sector. Namely, they address the high cost of downtime dwarfing the carrying costs of inventory resulting in "a 'just in case' approach to stocking spare parts inventory." This solution utilized a segmentation strategy to select four categories of inventory and generate inventory policies. By implementing the segmentation and review strategy, a net reduction of $47 million dollars was achieved during a period of service expansions. While the approach differs from our own approach, it provides a clear framework of the potential impact in terms of inventory reduction and provides avenues for additional improvements within the sponsor's inventory.

Schuh et al. (2015) approached the challenges of inventory management within the power industry from the perspective of a windfarm. Schuh et al. utilized a proportional hazards model, to examine core parts of the wind turbine and identify the risks associated with the equipment. While power poles do not have the same inspection cycle that a multi-million-dollar wind turbine has we still see this applied to the sponsor's approach. With regular inspections, areas of concern can be identified.

Once the need is recognized repairs can be planned and replacements can occur, ideally before a service interruption. With advancements in low-cost sensors and imaging technology associated with "Industry 4.0" predictive capabilities are likely to advance.

Two papers, Yoder (2013) and Kukreja et al. (2001) explore the reduction of inventory in the context of a network. Kukreja et al. follows Bailey and Helms as well as Schuh et al. in exploring the electric generation market. Kukreja et al. does so through the examination of 29 connected power generation facilities with common materials servicing multiple facilities. To support this review Kukreja et al. examines a single tier model with lateral cross-shipments between the facilities. This specific model provides a valuable example which matches the sponsor's current network. The sponsor operates in a single echelon with transshipments, making the Kukreja et al. framework the perfect model for their current state.

Yoder (2013) looks at a multi-echelon model with transshipments in the specific context of an electric or gas distribution network. The models Yoder examined consisted of central distribution centers supplying a broader network. Yoder's paper offers a basis which can be synthesized with the broader literature around multi-echelon frameworks. In developing his model Yoder points to insights from Mulvey et al. (1995) who proposed a robust optimization model. The proposed robust model would utilize a range of potential scenarios and find which solution best satisfies all of them. In the sponsor's view "the model is only as good as the next storm", which suggests that a robust approach considering performance in both storms and normal conditions is a necessity (Logistics & Support, 2021).

Yoder (2013) also paves the way to consider a broader range of literature which was not considered in depth in his review. He leveraged a model similar to that of Sherbrooke and the corresponding METRIC model. In synthesizing these two approaches we leverage the much broader literature around optimizing inventory within a network, into the specific context of an MRO inventory, with specific regard to an electric distribution network. This builds on the work by Yoder (2013) and Kukreja et al. (2001) and constitutes a unique contribution to the existing literature on power and utility companies.

2.4 Key Lessons

The METRIC Model, initiated by Sherbrooke (1966) has consistently improved. Garcia-Benito and Martin-Peña (2020) provide a summary of 36 different models around lateral shipments solutions across six different key criteria: identifying its focus on MRO, variable costs, transport time, lateral transport time, and model type. From this research we mapped a path for our own framework.

For consideration of demand Basten & van Houtum (2014) advise considering systems of line replaceable units (LRUs), not just single components. In effect, it would not matter if we had 90% of the parts to repair a downed power line, the line will stay down, and the customer will remain out of service. In support of this approach, Basten looks to Thonemann et al. (2002) who make a note that the systems approach, as opposed to a per part approach, provides significant benefit when costs are highly skewed.

Infrequent, but severely impactful events lead to the Robust Model proposed by Mulvey et al. (1995). Mulvey et al. (1995) suggests practitioners need to pay attention to periods of severe disruption for example after a wind or ice-storm where there may be both high demand, and longer transit times. For this we need to ensure we include both periods of regular operation and severe disruption.

For lateral shipments Kukreja et al (2001) found significant savings in enabling these shipments and so did Sherbrooke (1992). However, despite Sherbrooke's finding similar potential efficiencies, he also found that the lateral time needed to be one-quarter of the time of a base repair before it became effective. The sponsor will have to consider the increased complexity from a lateral shipment over that of a direct shipment and its impacts to lead time and cost. The sponsor has identified that lateral shipments can contribute to confusion and delay, and the operational complexity is a primary motivating factor in examining a new warehouse.

Combining the above learnings, we approach the development of first a network optimization model, optimizing the safety stock decisions according to stock levels and determining the impacts that network design has on final inventory parts.

3.0 Methodology

To compare a hub and spoke model and the sponsor's current decentralized network, we built a pair of network optimization models, handling both single echelon and multi-echelon planning, leveraging the lessons learned from the literature. We leveraged the sponsor's data from 2016 to 2021 to capture periods of high volatility, including severe weather events, as well as periods of normal operation.

Our methodology is divided into five sections:

1. Data Gathering
2. Model Structure
3. Demand Analysis
4. Validity Check
5. Detailed Model Construction

Through these steps we will create models which are realistic approximations of the real-world environment to provide an accurate assessment of their advantages and disadvantages for the sponsor.

3.1 Data Gathering

We received data from the sponsor detailing their operations since the start of 2016 to the end of 2021. This data included shipments between facilities, receipts, and consumption by the storerooms for both planned and unplanned activities. The data also included average lead times for products. We received the data regarding the breakdown of construction units (CUs), and outputs from their computerized maintenance management system (CMMS), and high-level details of the sponsor's logistics costs from one of their carriers.

In conversations with the sponsor, we identified that this six-year time frame will cover both normal operations in terms of line maintenance and expansion, as well as two major periods of disruption. The disruptions include both one of the most severe weather events faced by the sponsor in the past six years, and the wide-spread supply chain impacts from the COVID-19 Pandemic. In addition to these internal-facing events, the effects of several industry-wide disruptions including wildfires, hurricanes, and winter storms is present in the data.

We calculated descriptive data regarding the materials and inventory currently in use at the sponsor. Overall, the sponsor spends approximately $50 million annually, with a median monthly spend of $4.3m ranging between $2.9m and $5.1m. The sponsor completes work on approximately 1400 jobs per month; these jobs are typically small in value (median $249) but are skewed by a handful of jobs which are far more expensive. Showing the disparity in prices, 75% of all jobs completed cost less than $1,329 and 5% cost more than $11,600. This distribution is represented in Figure 1.

Figure 1: Histogram of spend by project

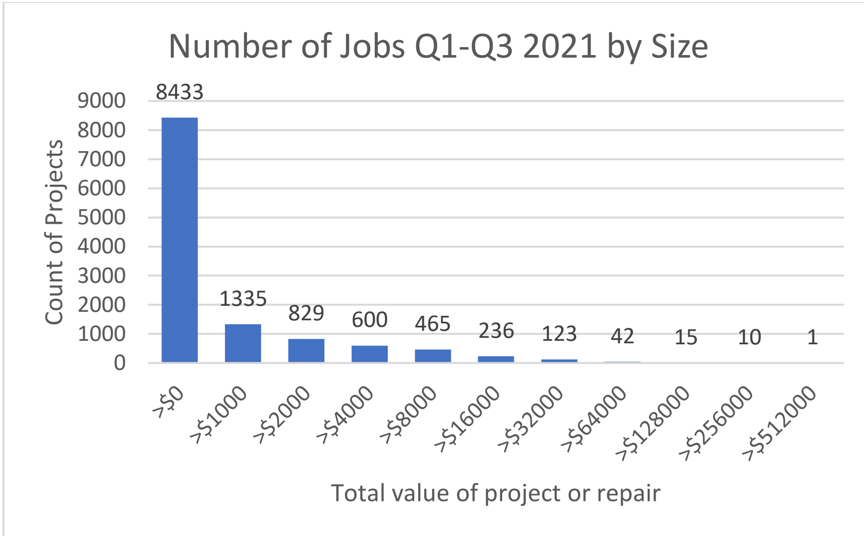

Based on the insights of Basten & van Houtom (2014), we examined whether the inventory is skewed. To do this we first consider materials on an as-used basis. The vendor has wire, for example, which is priced on a per foot or per pound basis but is often purchased or used in thousand-unit quantities. We leveraged the jobs to consolidate these into 'typical usage' quantities and found on a per item, per use basis, the data is heavily skewed (skewness 5.8) with significant outliers (kurtosis 24393). On this basis we know that a cluster of items will be important to understand the impact of inventory planning on service level.

We segmented the information into key items for our model to examine an average operation of the current network. The sponsor's data contained approximately 8,000 unique items and approximately 6.5 million records over the six-year period we examined.

Working with the sponsor, we refined the list to 2.5 million records relevant to electric distribution. We attempted to focus on materials which constituted significant demand and involved non-trivial prices. With these criteria we were able to further refine it to 1.3 million critical item records of demand. We then summarized this by month, this summary narrowed the list to 282k records of usage across the 2,895 unique items.

The data contained anomalies where materials were issued in one period and returned in another before being “reissued”. The net result of this would be a month of potentially negative demand where the storerooms would “grow” inventory from an uncontrolled source. To handle these anomalies, we zeroed 27k returns. This data cleansing results in the periods being treated as no demand. This helped maintain the conservation of flow assumptions.

3.2 Inventory Optimization Model Structure

We designed an inventory optimization model from the client’s current storeroom locations. Initially this is designed as a single echelon with transshipments. We leveraged a simplified list of 2,895 materials and aggregated these by month, the materials were selected based on their importance and spend. These two simplifications were based on advice received from Coupa Software and reflect limitations within the computational capacity available to us. After completing our analysis of the demand pattern, we further limited ourselves to 1,139 Stock Keeping Units (SKUs). We removed SKUs which had less than $5,000 total demand over the six years of data we had available. The remaining SKUs represent 28% of the total sponsor inventory and 77% of the relevant inventory.

The product data included actual lead times for materials from suppliers. In the event that no lead time was available for a SKU we utilized a 45 day lead time as a placeholder[5].

Key material information was also absent from our data set and was not able to be reliably generated. Product weights, stacking rules, sizes and minimum quantities were unavailable within the

---

[5] We compare this with a simple average of 31 days and a weighted average by price and quantity of 44 days. We chose 45 days as a simplified rounded number of the weighted average, specifically a month and a half, to represent a round number from the perspective of the business.

product master. We inquired with major vendors but were not able to find this information. This limited our analysis of logistics costs.

The absence of logistics information serves as an additional area for investigation in future research by the sponsor company. This gap also has two major impacts on the analysis. The first as previously mentioned is the absence of transportation savings from consolidation. However, for the creation of our model this limitation also resulted in significant simplification of the analysis. With no size information, aggregation of shipments from the central DC to the storerooms is no longer possible to model, as a result we aggregated shipments and demand monthly. Previously if this was done it would overstate the number of full truckload shipments under both scenarios. Second, we aggregated our supplier into a single point on this model. From a summary of purchase orders, the sponsor typically deals with 13 vendors who make up 87% of their spend. These vendors receive between 0.8 and 635 orders per month as shown in Figure 2.

Figure 2: Vendor distribution

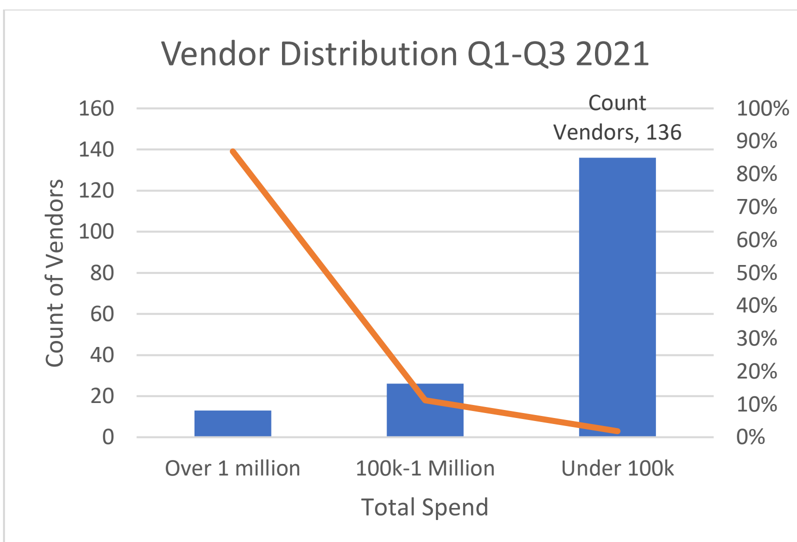

We used this opportunity to simplify the model to a single supplier and used an approximation of freight consolidation to better understand potential freight savings.

The next challenge was tracking the amount of demand which occurred within the lead time of the order. Reviewing the data from the sponsor's CMMS data we compared the date of an order's release to the ERP system against the log of the planning period associated with the orders. We found that jobs fell into two categories, planned and unplanned. Unplanned jobs represented 36% of the total demand. However, of the remaining 64%, the majority occurred within the lead time for the materials. The unplanned jobs are shown in the dashed areas of Figure 3, with the solid area in both charts representing material demand which could be planned for.

Figure 3 Demand by lead time

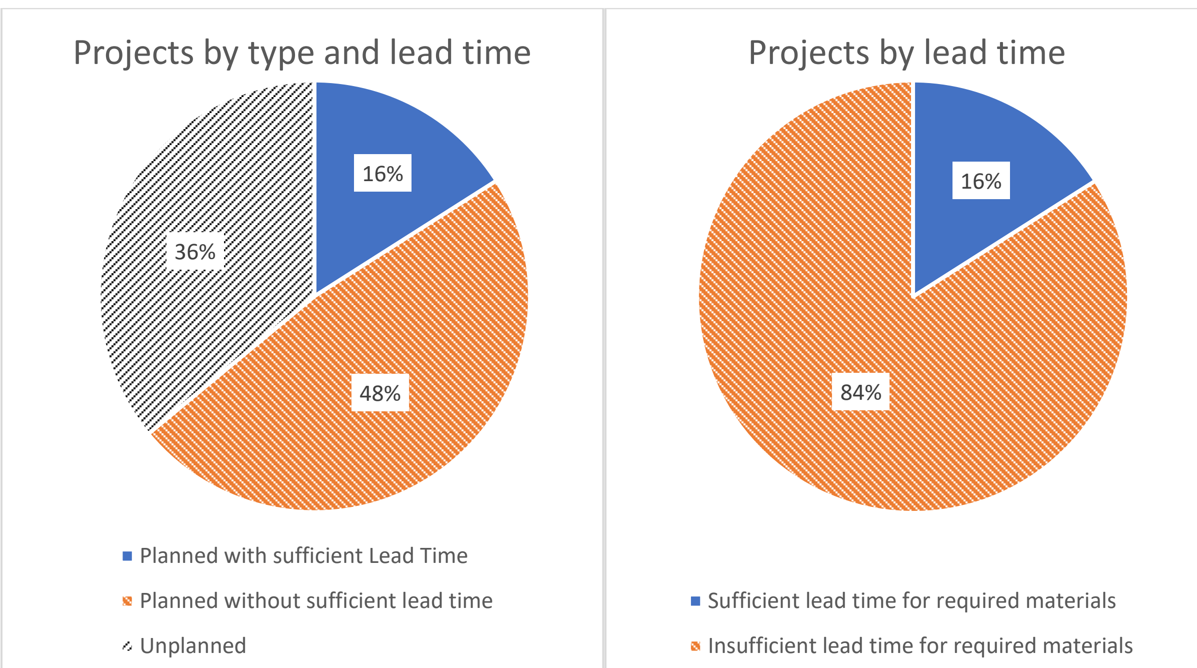

Together the unplanned and planned-within-lead-time demand represents 84% of all item usage. This strongly supports an order to stock methodology where numerous items for both planned and unplanned orders require an inventory. For our methodology we did not consider the 16% of demand which was forecasted with a sufficient lead time to be a substantial factor in safety stock planning, as a result we treated all demand as if it were unplanned.

In addition, we created a secondary assumption. In this case the storerooms were unable to source from the suppliers directly and instead sourced through the DC. We introduced the DC as being capable of shipping to any of the storerooms within two days. This estimate is based on the geographic area of the destination sites which are all approximately within three hours of a central point. This is expanded to two days to replicate the other challenges with delivery, including planning and scheduling. We view these to be fixed across all shipments and the primary determining factor of lead time.

3.3 Demand Analysis

To classify demand, we examined the sponsor's historical demand data from 2016 through 2021 and leveraged Supply Chain Guru's built-in demand analysis. This allowed for a classification of demand as either intermittent or non-intermittent.

Non-intermittent demand is characterized by near continuous patterns of usage with a few sparse gaps in the data. This can further be refined into two subtypes of demand: erratic and smooth. Erratic demand is categorized by one or more periods of high variability within the review period, whereas smooth demand is shown to have few periods of no demand within the review period. Additionally, smooth demand seems to display characteristics of level, trend, and potentially seasonality.

Conversely, intermittent demand often has periods of no demand for two consecutive periods or more. Intermittent patterns are also defined by high variability and low variability. They can then be further divided into slow and lumpy usage patterns; where slow demand is often characterized by large gaps of usage and lumpy demand is followed by intermittent periods of high usage, then little or no usage. This classification is summarized in Figure 4.

Figure 4: Demand classification flow chart (Coupa, 2021)

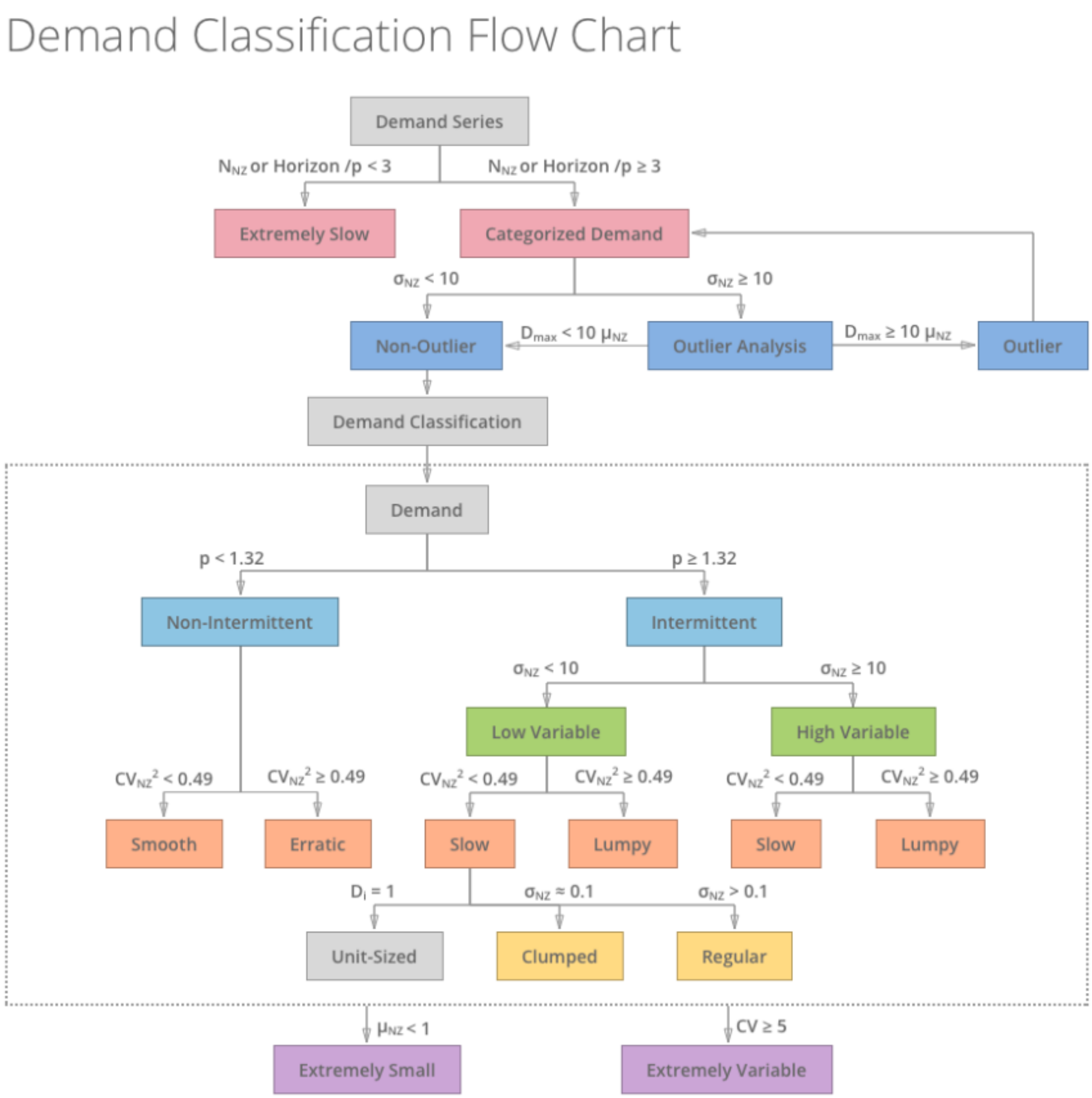

One transformer's demand across 15 different sites is categorized in multiple ways depending on the site, but when aggregated, the demand normalized into a less variable, non-intermittent, smooth pattern shown in the five charts of Figure 5. Smooth distribution is the ideal pattern since it lends itself to more reliable forecasting.

Figure 5: Demand summary across categories

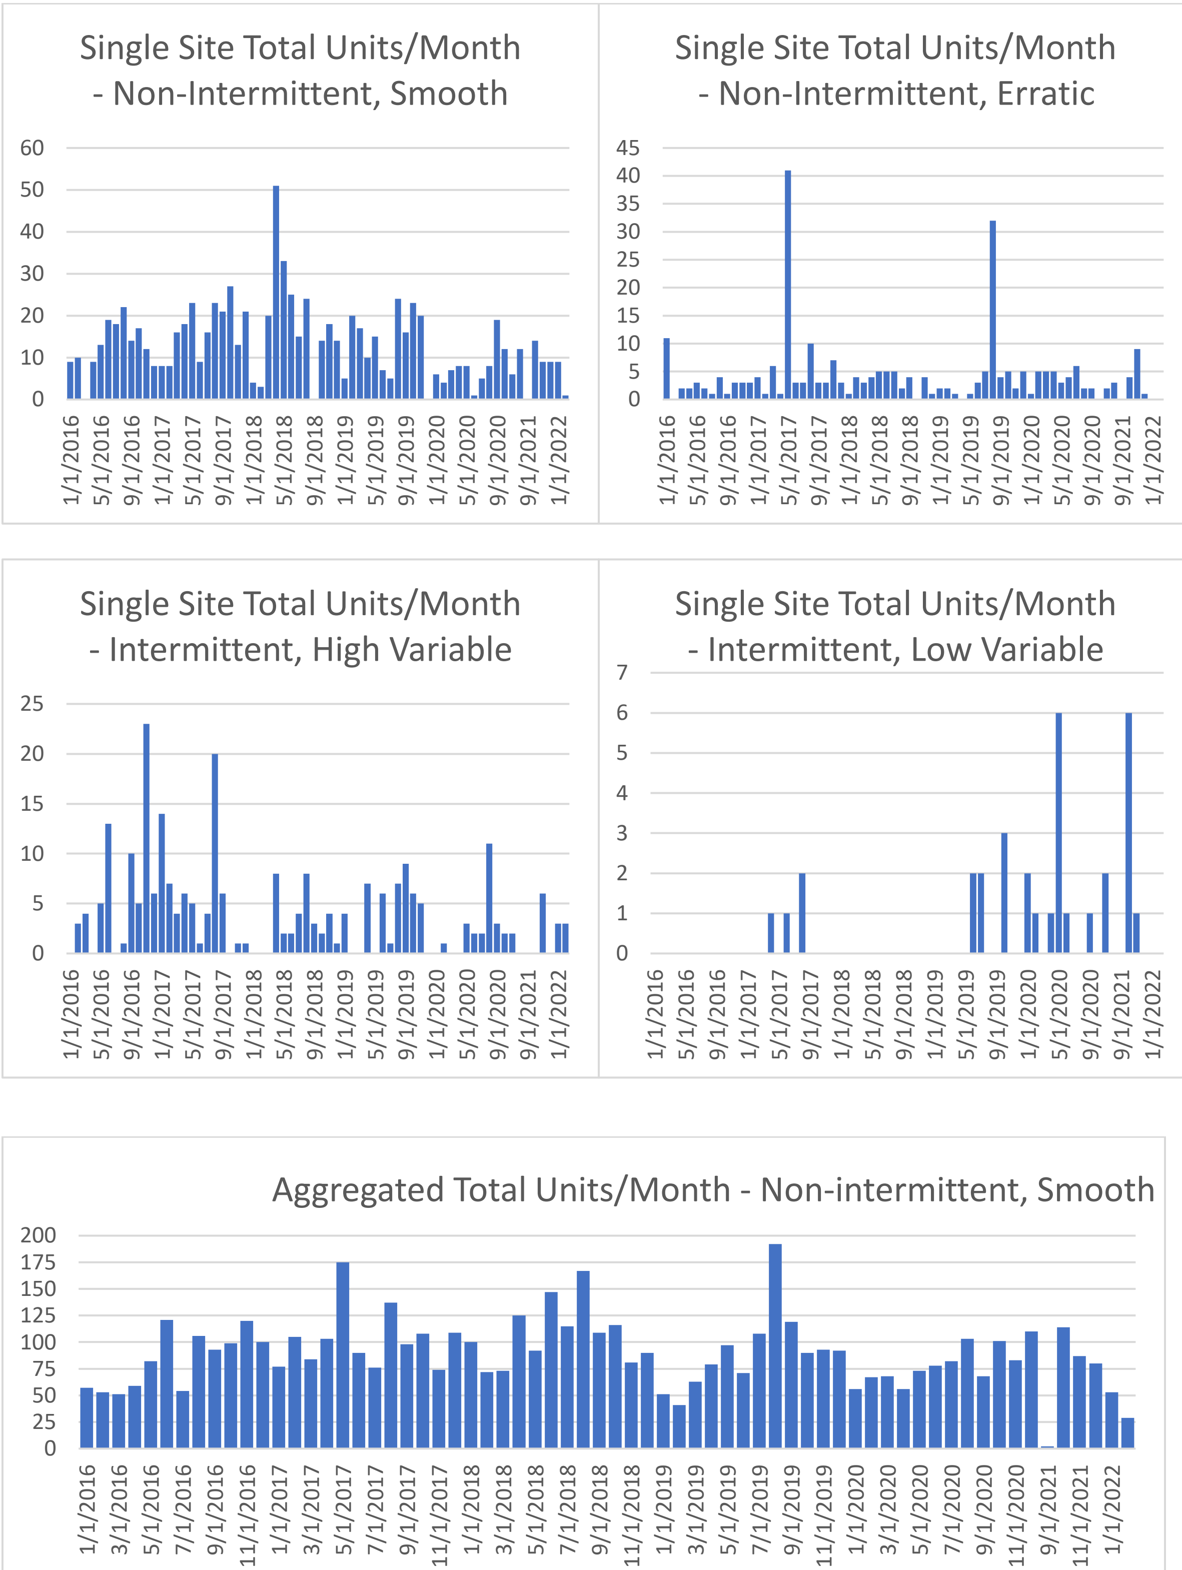

3.4 Validity Check

We first completed a demand analysis of both a single DC model and the existing decentralized model to determine if sufficient savings were feasible. We then classified the demand into buckets as shown in Table 1 using the standard classification employed by Supply Chain Guru as described in

Figure 4. As expected, the impacts of aggregation resulted in decreased counts of distributions across all descriptions except smooth demand:

Table 1: Demand classifications in centralized and decentralized models

Demand classification	Current storeroom decentralized demand patterns	Central DC patterns	Reduction
Erratic	693	296	57%
Extremely Slow	5841	526	91%
Extremely Variable	1936	259	87%
Lumpy	5515	926	83%
Slow	4369	565	87%
Smooth	256	263	-3%
Grand Total	18610	2835	85%

In this analysis, as would be expected if the patterns were stable between the DCs, centralization does not simply reduce the overall number of SKU patterns. Instead, we see the smooth distributions increase, with outsized impacts in the reductions to the Extremely Slow, Slow, and Extremely Variable patterns. With patterns that are similar or highly dependent on each other, we would expect a distribution like the one identified in Table 2. The expected patterns column displays our calculation multiplying the percentage of occurrences in the decentralized model of Table 1 by the reduced number of SKUs[6].

[6] We find the results under this analysis to be significant based on a Pearson Chi Square test to .05 P-Value (alpha <0.00001)

Table 2: Counterfactual, a simple reduction

Demand Classification	Calculated Patterns	Expected Patterns
Erratic	296	106
Extremely Slow	526	890
Extremely Variable	259	295
Lumpy	926	840
Slow	565	665
Smooth	263	39
Grand Total	2835	2835

To give a comparison of the potential impact we examined the standard deviation over the lead time for the demand analysis both for decentralized inventory as well as centralized inventory. We note that risk pooling reduces the overall inventory by approximately 54%, as shown in Table 3.

Table 3: Demand over lead time for storerooms under decentralized and centralized inventory

Demand Classification	Decentralized Inventory	Centralized Inventory	
	Cost 1 SD Inventory	Cost 1 SD Inventory	Reduction
Erratic	$828,059.73	$1,085,325.54	-31%
Extremely Slow	$1,609,110.59	$139,829.44	91%
Extremely Variable	$1,284,179.71	$250,628.49	80%
Lumpy	$6,028,163.86	$1,570,493.79	74%
Slow	$3,778,918.29	$1,141,823.69	70%
Smooth	$505,630.75	$2,327,579.73	-360%
Grand Total	$14,034,062.93	$6,515,680.68	54%

This distribution suggests improved performance from an inventory perspective. The central DC benefits from this decreased variability as it decreases these difficult to plan for SKUs over the lead time of restocking inventory. A reader might expect that in a reduction of storerooms from 72 to 1 (98.6% reduction) that there would be a higher impact than a 54% reduction in safety stocks. However, not every part is present in every storeroom meaning that not all SKUs experience a full reduction through consolidation. The impacts of this decrease in complexity extend beyond the simple reduction and in section 4.2.2 we discuss the broader implications of the complexity reduction.

The significant difference in safety stocks between these two scenarios overstates the overall impact. The sponsor has previously worked to improve performance through the practice of transshipments, similar to what was discussed by Kukreja et al (2001). In this process a storeroom does not need to rely solely on their own stocks, but the inventories of all the storerooms can collectively serve to provide the safety stock to handle the volatility at any given one of them.

However, the sponsor reported difficulty administering transshipments. Three areas of delay exist. These include the difficulty of organizing lateral shipments, difficulty in determining the amount of unfulfilled demand in the facilities, and delays with sending small less than truckload (LTL) shipments. We discussed these impacts in greater detail in section 4.2.1, however, for our purposes we estimate a two-day lead time for a direct DC arrangement and a seven-day lead time for a lateral shipment. With these assumptions, we compare where a site needs to only carry inventory to handle the variable demand over a lead time of two days for the shipment from a DC or seven days from the shipment from another site. This difference is shown below in Table 4.

Table 4: Changes in storeroom safety stocks from faster replenishment associated with a hub replenishment

Demand Classification	1 SD over 7 Days	1 SD over 2 Days	Difference
Erratic	$333,031.60	$178,012.88	$155,018.72
Extremely Slow	$522,269.80	$279,164.95	$243,104.85
Extremely Variable	$393,014.26	$210,074.96	$182,939.30
Lumpy	$2,102,874.14	$1,124,033.51	$978,840.63
Slow	$1,119,774.88	$598,544.85	$521,230.03
Smooth	$159,557.52	$85,287.08	$74,270.44
Grand Total	$4,630,522.20	$2,475,118.23	$2,155,403.97

Using this model, we build a broad estimate of the impact that we expect in a full multi-echelon strategy. By moving to a DC, we expect to pool the inventory across the order lead time, while reducing the amount of inventory for sites which are not covering the order lead time. From here we compare the likely expenditure of a new warehouse against the expected savings in terms of capital to determine the high-level feasibility of this approach. From this initial data analysis, we conclude that consolidation offers significant potential savings.

3.5 Detailed Model Construction

We designed two models to test our hypothesis. Effectively each storeroom should hold inventory for the items they are likely to use to cover the expected lead time based on their total service level. For a single storeroom the equation for optimized inventory iterated per item is shown in (1). This results in a calculated optimal inventory for a comparison point to the hub and spoke model.

Under the hub and spoke model decision variables are introduced to determine whether the inventory is held at the DC or in the storeroom as shown in (3). These decision variables are then subject to the constraint that the inventory levels for each product and location must both be equal to or higher than the variability of demand for the lead time relative to the cycle service level as shown in (4) and (5).

Safety stock formulation decentralized model

$$ss_{ij} = Z_{\alpha} \sqrt{E(L)_i \, \sigma^2_{D_{ij}}} \tag{1}$$

$$ss_{\$\,total} = \sum_{j\epsilon J} \sum_{i\epsilon I} ss_{ij} p_i \tag{2}$$

$ss_{\$\,total}$	Total Safety Stock in dollars
ss_{ij}	Safety stock target inventory level for material i in storeroom j
Z_A	the inverse of normal cumulative distribution of service time
$E(L_i)$	the expected lead time by item
$\sigma^2_{D_{ij}}$	the standard deviation of demand by item by storeroom
ss_{ij}	the safety stock per item per storeroom
p_i	the price per unit

Objective variables for hub and spoke model

$$\min\left(\sum_{i\epsilon I} p_i \left(X_{iq} + \sum_{j\epsilon J} X_{ij}\right)\right) \tag{3}$$

Service level constraint for a single DC

$$\sum_{i\epsilon I} \frac{X_{iq}}{\sqrt{E\left(L_{iq}\right) \sigma_{D_{iq}}^{2}}} \geq Z_{\alpha} \tag{4}$$

Storeroom service level constraint

$$\sum_{j\epsilon J} \sum_{i\epsilon I} \frac{X_{ij}}{\sqrt{E\left(L_{ij}\right) \sigma_{D_{ij}}^{2}}} \geq Z_{\alpha} \tag{5}$$

Z_α	the inverse of normal cumulative distribution of service time
$E(L_{iq})$	the expected lead time by item for the DC
$E(L_{ij})$	the expected lead time by item for the storeroom
$\sigma^2_{D_{iq}}$	the standard deviation of demand by item for the DC
$\sigma^2_{D_{ij}}$	the standard deviation of the demand by item for the storeroom
p_i	the price per unit
D_{ij}	the standard deviation of demand by item for the storeroom
D_{iq}	the standard deviation of demand by item for the DC
X_{iq}	the decision variable for stock held in the DC
X_{ij}	the decision variable for stock held in the storeroom

In the single echelon design, shown in Figure 6, each location calculates its inventory to cover demand over lead time. Transshipments are allowed, but not planned for. This construction replicates the current practice using transshipments to spread out economic order quantities, but each storeroom plans its safety stock independently (Logistics & Support, 2021).

Figure 6: Decentralized Supply Construction

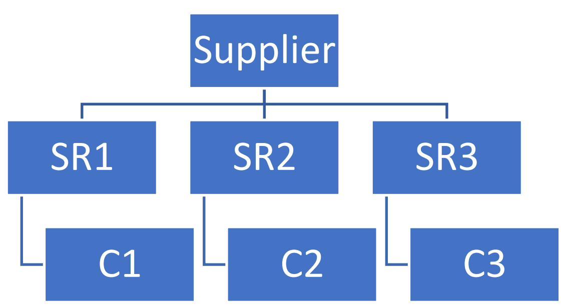

By contrast, in the hub and spoke model the system calculates an expectation of stocking out either over the lead time between the storeroom and DC, or between the DC and the Supplier, reflecting the full coverage period for the network as shown in Figure 7. Beyond the difference of the insertion of the DC into the center of the chain and the requirement that the storerooms source direct from the DC instead of direct from the supplier, the models are otherwise the same.

Figure 7 Hub and spoke framework

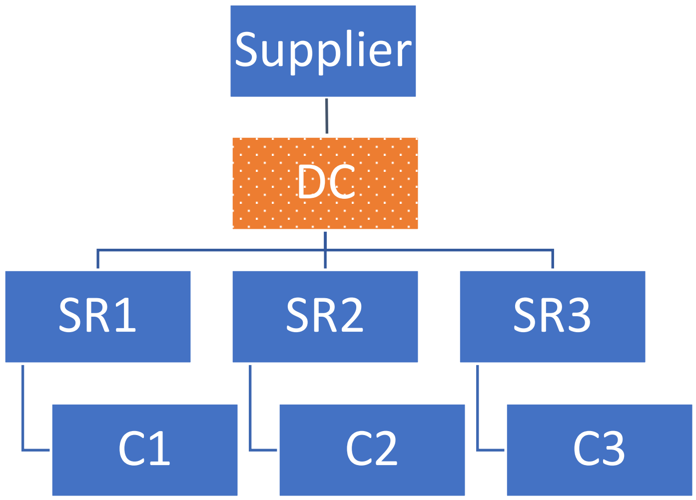

We then iterated across different cycle service levels using 90%, 95% and 99%. The details of the model inputs are discussed in section 3.5.1

3.5.1 Model Inputs

3.5.1.1 Storerooms

We modelled each storeroom as a unique location with unlimited capacity for parts sourced directly from either the supplier in the decentralized model, or from the DC in the hub and spoke model. In the analysis, we identified that some of the storerooms listed within the data were not primarily serving the electrical grid and were involved in pulling smaller amounts of common materials between the sponsor's business units. Under the assumption that the sponsor would likely not source these storerooms under a separate policy these were left in. This design would reflect the idea that for a category of items sourcing would transition for all sites. For the sponsor, there exists an opportunity to improve the computational efficiency of the model by removing these storerooms in future iterations.

3.5.1.2 Customers

We fed demand into the storerooms from the customers with a one-to-one relationship between storerooms and customers. We viewed the customers demand to be satisfied when the storeroom kitted the product for the workers. We provided no lead time between the customer and the storeroom. Customers in this sense reflected both explicit work orders from a customer, and the general maintenance requirement for the area of the grid being serviced by a particular storeroom.

3.5.1.3 Suppliers

As discussed in section 3.2, the absence of logistics information allowed us to model suppliers as a single source. We set up the supplier as having unlimited capacity and instantaneous production times. Instead, sourcing lead times were supplied within the transportation segment to ensure that the model would hold inventory in the DC or storeroom but not at the supplier. This choice is an approximation from the perspective of the sponsor, due the terms of their long-term contracts. It was a common perspective, including in METRIC models, such as the formulation by Sherbrooke (1966). The supply chain disruptions from the COVID-19 pandemic have identified limitations with this approximation and offers an opportunity for further development in future research by the sponsor.

3.5.1.4 Holding Cost

The sponsor identifies a range of costs similar to a holding cost on a per storeroom basis for their internal metrics and regulated return functions. For our modelling, we utilized a 25% holding cost across all storerooms. The constant holding cost does not impact the optimization model. The internal allocations are also primarily fixed costs, when combined with the fact the sponsor was not looking at closing or opening storerooms, this aligns with the sponsor’s intent (Logistics & Support, 2021). Should the sponsor choose to examine the number of storerooms the specific details of warehouse costs should be more fully modelled in that future research.

3.5.1.5 Transportation Lanes

Transportation lanes were modelled from the supplier to the DC as zero days, compressing the lead time into a single lead time which was part of the sourcing lead time. This assumption makes no distinction between the portion of the lead time spent on the truck, versus the lead time being sourced.

3.5.1.6 Materials & Lead Times

As discussed in section 3.2, we focused on a limited number of 1,139 SKUs. The product data from the sponsor included actual lead times for materials recently purchased from suppliers. For items if no recent lead time was available, we utilized a 45-day lead time based on the weighted average lead time. Under the decentralized model, these lead times were experienced directly by the storerooms. Under the hub and spoke model, the DC experienced the lead times from the supplier, and the storerooms experienced an additional a 2-day lead time from the DC as shown in Figure 8

Figure 8 Lead Time implementation in hub and spoke and decentralized models

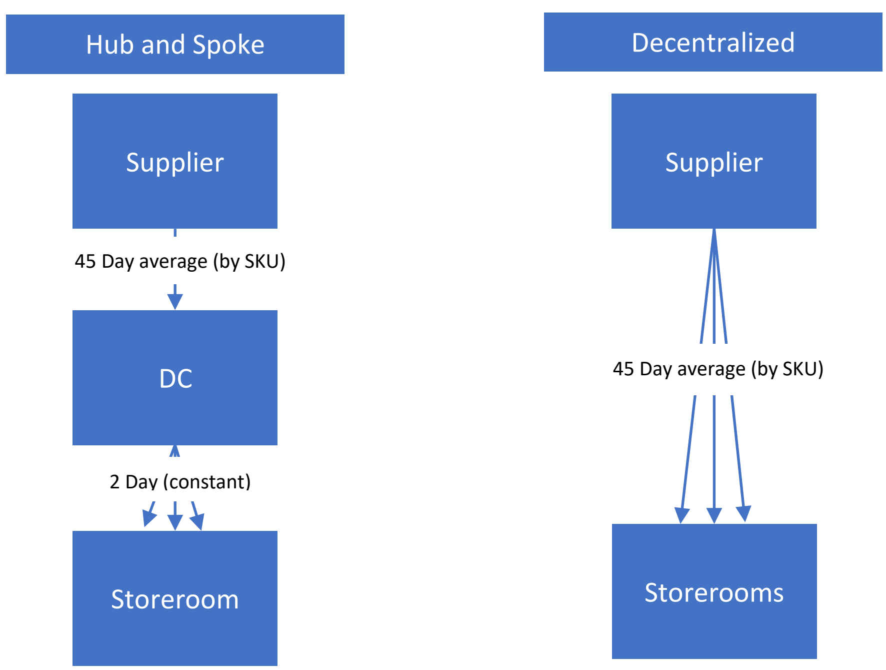

3.5.1.7 Demand patterns

We used the historical material consumption from each storeroom as a proxy of demand. This method has the advantage that it creates a clear indication of when demand occurred and how much was used. This method also has the disadvantage that it does not show unfulfilled demand. This model can be improved with future measures of unfulfilled demand as they become available.

3.5.2 Cycle Service Level vs Item Fill Rate

For this paper, we primarily focused on cycle service level. We believe that missing one part of a job means that the job cannot proceed. The total sum of parts is required for the job to be completed in full. This design is distinct from the Item Fill Rate, which may be more appropriate in a consumer goods segment where each unit might be separable for the quantity demanded. We believe that cycle service level is the appropriate criteria for the sponsor's case because a requirement for electrical wire would likely require the entire span and being able to partially fulfill the span would be insufficient. A further improvement to this model would be to consider Line Replaceable Units (LRUs), as discussed in section 4.2.3.

With these inputs cleaned and set, we built our optimization model and completed six optimizations forming the basis of our analysis.

4.0 Results and Analysis

We found that under a high service level condition, a hub and spoke model will likely lower overall inventory levels while improving service levels.

A single dataset was used, consisting of the full period of demand from 2016-2021 and for both network designs. The first model consisted of a single echelon planning for the storerooms on their own, with the ability to resupply from each other, but without planning on that basis. The second model allowed for a multi-echelon strategy where inventory could be consolidated at a central DC. Results from both models are shown in Figure 9.

Figure 9: Impact of network design on required safety stock

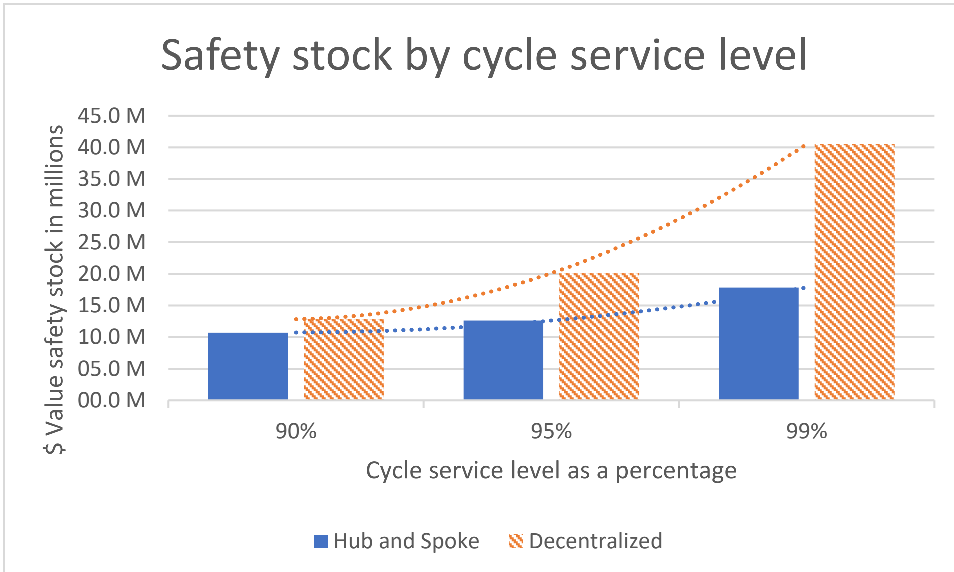

Under a low service condition (target service level 90%) the addition of the DC sees nearly equivalent inventory levels. The benefits from a centralized warehouse would be unlikely to pay for the cost of the warehouse. However, as the target service level increases a decentralized model more rapidly increases its safety stocks. This result matches intuition. At a lower level of inventory, the additional inventory from establishing in a central DC would counterbalance the savings. However, because the central DC is exposed to less volatility, increased service targets can be achieved with lower inventory levels. Notably, with a multi-echelon model a 99% service level can be achieved with slightly less inventory than a single echelon model is able to deliver at a 95% service level.

With a 98% cycle service level, the decentralized model requires a safety stock of approximately $31.4 million dollars. Combined with the sponsor's average cycle stock of $5.5 million dollars, the total inventory compares closely with the $38.6 million dollars observed for 2020 year-end for inventory for the relevant stocks. This suggests that over the long term an individual item cycle service level of between 98% and 99% can be obtained with inventory at current levels.

4.1 Baseline Comparison

To determine validity, we will examine the service level and the baseline inventory policies of the sponsor against the model's calculated inventory stocks.

A traditional approach to determining inventory levels relies on the analysis of service levels and determining the likelihood of a stock out. To determine the extent to which the model reflects the sponsor's experiences it is important to understand the current service levels experienced by the sponsor. A review of their data suggests that the reliability of the grid is exceptional, with a reported grid uptime of 99.9%. During a severe weather event they were also able to restore 95% of their customers within six days. However, a review of the planned projects suggests that only 75% of the jobs are completed within 30 days of their target date.

The reason for the discrepancy in the observed metrics speaks to the complexity of fulfilling demand. A 98% item service level would not result in 98% service level for the individual construction units (CU). A CU contains between 1 and 47 materials with an average of 5. Across all CUs we would expect a 98% service level at a material level to result in a 92% service level at the CU level. Projects and repair jobs are themselves often aggregations of multiple CUs, resulting in increased complexity and required parts, with a corresponding decrease in service level. Service level targets are discussed in greater detail in section 4.2.3.

We observe that the minimum and maximum stocks add up to $14 and $20 million dollars respectively, while actual inventory adds up to $38.5 million dollars. Calculating the cycle stock we would expect an average of $5-$6 million dollars of inventory based on the average ordering patterns of the various sites. This represents a large discrepancy between the targeted and actual inventory levels.

Without a clear understanding of the current service level experienced and the extent to which inventory drives the result, it is impossible to determine the extent to which inventory is being buffered above the set inventory levels by better-than-expected forecasting, unofficially elevated safety stock levels, or simply inefficiency in executing to a safety stock from demand that evaporated or product obsolescence. Further, we know that an optimized model will always portray a lower inventory and higher service level than actuals as it is functionally very difficult to maintain optimality.

For our analysis we are interested in the impacts that a central warehouse might have, regardless of the specific method inventory policies set within the facility. In this light, the service level currently achieved by the sponsor is less interesting than what the target service level should be. To answer this question, we consider the impacts of independent cycle service levels across multiple items, which are required for the same order as shown in

Figure 10.

Figure 10 Mathematical relationship between cycle service level and full order rate assuming independent items

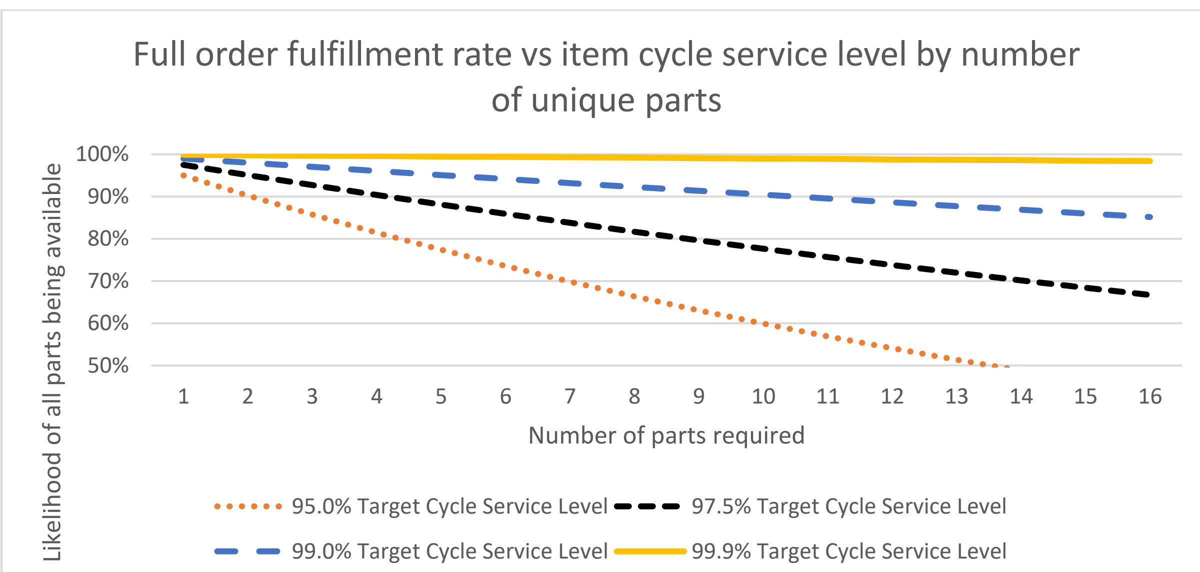

With this viewpoint we see that a high item cycle service level is required to maintain even a 75% cycle service level for a job with 10 distinct parts[7]. On this basis we recommend that an appropriate target inventory on an average basis would be above 95%. This threshold is of particular interest to us because this is the region shown in Figure 9 where a hub and spoke model begins to save sufficient money over that of a decentralized model. Based on the current inventory and the implied service level that we recommend targeting a 99% cycle service level.

We believe that our models provide a reasonable approximation of the potential savings. The specific amount of savings to be provided versus the expected service level should also be considered in terms of the non-modelled aspects discussed in section 4.2.

4.2 Impacts not modeled

There are several impacts that are known but were not captured in the data we examined and therefore were not modeled and measured; however, these factors represent potential savings to the sponsor. The benefits from volume discounts through order aggregation, reduced shipping costs because of full truckload consolidation on both inbound and outbound loads, reduced process complexity on ordering, receiving, and return processes, as well as increased cash discounts are all realizable. We have summarized the major impacts that were not modelled in Table 5.

---

[7] There are two major types of demand, planned work and unplanned work. Aggregating the data by day and storekeeper we identified that for unplanned demand a median of 10 unique part numbers were pulled each day with an average of 13.

Table 5: Summary of impacts not modelled

Impacts not modelled and impact on benefits of a hub and spoke model	
Improves benefits	Decreases benefits
Order consolidation lowering freight costs	Including forecasted demand
Lower cost of ordering	Faster transshipments
Decreased process complexity	
Simplified reverse logistics	
Volume Discounts	

4.2.1 Full Truckload considerations & Costs of Ordering

As mentioned in section 3, we were unable to develop sufficient data to model the impact of a central warehouse on freight costs. Any impact would be beneficial and not substantially change the end conclusion. However, it is important to consider the impacts of the freight costs. Reviewing the initial data set between 2021-01-01 and 2021-09-30 it was identified that there were 22,660 unique purchase orders placed across the network, in addition to these orders direct from vendors there were 785 unique transshipments executed over the same period, amounting to a total volume of replenishments of 23,445 restock orders, either from other warehouses or from vendors.

If we assume that each of the 72 storerooms receive shipments every weekday and utilize the unique week/vendor combinations to assume that DC replenishment occurs at most weekly, this works out to an estimated 21,800 shipments. This is less than the number of shipments currently being created, and we would expect a higher service level. If we use the unique number of creation dates per storeroom as an estimate, we find that 7,086 orders would be necessary in this period. This represents a reduction in complexity for the sponsor of roughly 70%. This complexity reduction is a cost from reducing management time spent on processing, planning, and executing these orders. Assuming 15 minutes per order and a burdened labor rate of $25/hour equates to $102,000. However, we do not believe it would be beneficial to achieve these cost savings from headcount reduction. Instead, we recommend the sponsor look at how these workers could support other initiatives within the organization. Based on our observations there are significant opportunities that workers could address, ranging from improving the quality and extent of the data available within the sponsor's ERP system, to initiatives to improve forecasting and customer outreach which has been identified by the sponsor as an ongoing challenge. These savings are unique in that they help serve to decrease the inefficiency costs within the network. Some common areas of investment include better determining the unfulfilled demand rates for materials, and the specific causes of delay within projects for future improvements. These presented challenges for this project as identified in section 4.1, but they also pose challenges for future process improvements. By freeing up resources the sponsor could accelerate process improvements within the organization.

4.2.2 Process complexity & Volume discounts

The impact to logistics and ordering costs are aspects of the complexity caused by a decentralized system. Similar to the analysis laid out in section 4.2.1 a decentralized network results in additional burdens on the ordering team than a hub and spoke model imposes. It is useful to consider the opportunity cost associated with what those employees could be doing otherwise. The sponsor has identified to areas of priority, the first is pursuing volume-based discounts and the second is accessing payment discounts on offer from suppliers (Logistics & Support, 2022). For example, an annual category spend of around $50 million dollars on parts with an incremental 1% discount represents a savings of $500k. This presents the opportunity to further offset the cost of the warehouse.

4.2.3 Line Replaceable Units vs Item Cycle Service Level

As briefly discussed in the literature review, Basten & van Houtum (2014) recommend considering items as line replaceable units (LRU), particularly when item values are highly skewed (Thonemann, Brown, & Hausman, 2002). Based on our data analysis, the sponsor's item values are highly skewed, however, as identified in section 3, we were unable to include LRUs and instead relied upon the impacts of individual cycle service level. A per item service level will always be higher than the ability to fill the full requirements of a maintenance order, and jobs might consist of multiple CUs each containing multiple materials.

By way of example, consider the replacement of a single power pole in a residential environment. We have CUs for the foundation, pole, the span of wires, and for the cross arm. This relatively simple job encompasses multiple CUs and each CU consists of multiple parts. The high number of distinct parts suggests that even relatively high cycle service levels can result in much lower actual performance rates. This is partially ameliorated if an alternate material or BOM can substitute for one which cannot be filled. For example, a stronger cross arm could substitute for a weaker one.

For the sponsor, a solution is to simply hold a higher service level across all parts until such time as the data supports a more detailed analysis to optimize parts on an LRU basis. Our analysis supports the sponsor targeting a cycle service level in the high 90th percentile. We would recommend targeting at or above 99% as it is close to what is currently possible to achieve with existing inventory levels.

Assuming a 25% cost of inventory and a $1 million budget for a warehouse, the beneficial impact of a hub and spoke model was found above a 95% service level. A 95% service level would reflect the current min-max order levels that the sponsor has, and a 98-99% service level would reflect the current inventory that the sponsor has of the target materials. We conclude that, despite the limitations of the analysis, a hub and spoke model would pay for itself. We also recommend, because of the impacts of cycle service level on LRUs, that the sponsor substitute a high cycle service level as their target inventory metric. Process improvements can provide better optimization, as covered in the literature review, advantages can be had by ensuring that the most expensive items in an LRU are carefully managed, while maintaining very high service levels for more affordable parts. The primary barrier to completing such an analysis is the time required to identify and document the primary LRUs as well as substitutions which can be made. In this regard implementing a hub and spoke model will provide initial savings and provide the time savings necessary to allow the investigation of how best to document LRU which would allow a better optimization of inventory in line with the literature review.

5.0 Discussion

From the results and limitations laid out in section 4, we identify multiple potential areas of savings from a hub and spoke model, assuming a 99% service level we expect savings according to Table 6. We would not expect that all identified savings would be realized in the first year. However, we believe that even with conservative estimates for both warehouse cost and initial savings that the warehouse would provide benefits which would exceed its costs in the first year.

Table 6: Expected Savings at 99% cycle service level

Category	Approximate Savings
Holding Cost of Inventory Reduction @ 25%	$5,700,000
Reduced Ordering Costs	$100,000
Bulk Discounts	$500,000
Logistics Costs	Out of Scope
Total Potential Savings Estimated	$6,300,000

We identify two alternate ways that the sponsor could achieve the reductions identified in this analysis. The comparative difficulties and costs associated will inform the best approach. The factors are demand volatility and lateral shipment time. A hub and spoke model addresses both issues but is not the only method to do so.

Demand volatility can be addressed through improved forecasting, both involving estimations of likely causes of failure and explicit service constraints. Decreased volatility would drive lower safety stocks. Just as the sponsor could decrease their volatility by pooling, it could decrease the volatility through better forecasting. The challenge here is determining the amount of forecasting that the sponsor could realistically achieve. Natural disasters can be predicted insofar as they can be realized to be a risk, but it is difficult to forecast them to a time and date several months out. More regular maintenance can be predicted, particularly with improved sensors and observation, but not all parts which fail will fail because of predictable wear and tear.

Lateral shipment can also be improved between sites and the sponsor could look to closely manage replenishment between sites as an alternative strategy. This would mirror the approach utilized by Kukreja (2001). We recommend against this approach. We expect that as the already discovered by the sponsor, administering lateral shipments is significantly more complicated and this burden would slow other improvements that the sponsor is seeking to implement (Logistics & Support, 2021). Further, as found by Sherbrooke (1992) this is primarily useful when the lateral shipment would be significantly faster than a resupply from a hub. Based on the sponsor's geography we do not expect a delivery from one warehouse to another to be one quarter of the time once administrative and scheduling delays are taken into account.

5.1 Implementation Strategy

Next, we map the recommendations for implementation of the insights garnered from the research results and discussion outlined above. The roadmap encompasses the costs, structure, risks, and opportunities. However, it is not our intent to enumerate every potential risk and benefit to the sponsor, but rather highlight the most immediate and likely.

There are several factors beyond the cost of the facility to consider when implementing a hub and spoke model. Change management will be crucial for the sponsor's successful implementation and costs, such as labor to staff the warehouse, equipment and infrastructure to operate, the inventory itself, and the logistics cost of the new plan will need to be considered. Labor agreements are already in place, so any changes would be subject to the provisions of the agreement and the deciding factors would be limited to staffing levels required to attain the desired service levels. Likewise, long-term contracts already exist for equipment and facilities infrastructure for the storage needs of the new facility. These could be leveraged to ensure the costs of the proposed central model remain sustainable.

The first risk to the successful implementation is the overall lack of a structured framework surrounding and supporting the ability of the sponsor to accurately measure and report the desired level of metrics, such as the level of unfulfilled demand, the construction unit service level, and logistics efficiency. The absence of these metrics makes justifying the hub and spoke model substantially more difficult. Establishing this structure should be a top priority before further steps are taken.

An additional risk will be getting and maintaining stakeholder buy-in to overcome the siloed nature of the current operations. Without this buy-in, these silos will become caches of inventory as operations teams hoard inventory to offset perceived disruptions from the central warehouse. There is also the question of reverse logistics. Implementation will require careful consideration to avoid the previously mentioned caches. The central warehouse must accept returns and process them to prevent inventory levels from growing needlessly at the spokes.

The opportunities are multi-dimensional and implementation dependent, with varying degrees of benefit recouped. First, optimized inbound and outbound material flows provide a more efficient movement of materials and goods into and throughout the network. One further benefit of these improved flows is capturing the time savings at the spokes relating to inbound receipts and the need for reduced lateral shipments. Further, this presents the opportunity to establish a more efficient reverse logistics process, providing for a more expedient return of materials, resulting in more time savings.

Additionally, there are process improvement opportunities around improving the internal inventory audit processes, the receiving and payment processes, as well as the data collection and maintenance of materials from a systems perspective. The lack of this data on material characteristics was a limiting factor to our analysis and we excluded logistics evaluations because the assumptions we would have to make would be too broad and general to be relevant or applicable. However, we believe investing in these process improvements will lead to both time and cost savings that may be in addition to other savings proposed. Still, these values are not explicitly quantifiable at this state and the change management will continue to be crucial before, throughout, and following the iterative implementation.

5.2 Storm Stocks

A limitation on this model is that the aggregation of demand is largely independent. One of the areas of concern initially identified by the sponsor is that an inventory model is only as good as the next storm (Logistics & Support, 2021). As noted in section 3.1, we included periods of significant disruption in our dataset to ensure we capture these disruptions in the variability of demand. However, this model might not fully reflect the dependency of demand between storerooms. An area of focus for the sponsor is improving its storm stock inventory while simultaneously decreasing the general inventory it is carrying (Logistics & Support, 2022). We still recommend the hub and spoke model in this event as it allows the sponsor to minimize inventory generally and to focus on storm stocks as a separate category. We believe that storm stocks could be composed of a small number of items where more expensive items might substitute for cheaper ones in order to simplify logistics during a major disruption.

6.0 Conclusion

Based on the analysis of the potential inventory reductions and the ability to improve service level at a reduced cost we recommend that the sponsor implement a hub and spoke model to reduce the cost of inventory. We recommend that the expected time reduction for the teams in administering this network can be reinvested into improving the information available and continuing investigations into the overall network design, as it is our belief that this represents only part of potential cost reductions. In particular, we believe that implementing a hub and spoke model can both provide immediate cost savings in this context and also provide a mechanism to free up management time to reinvest into additional cost savings.

As a broader expansion, we would expect that the hub and spoke model be considered for any capital-intensive business with multiple storerooms spread over a moderately sized geographic region. We believe these findings are relevant not only to power utilities but also other utilities including water and pipelines. We encourage any researchers or professionals examining our analysis to consider the application of a hub and spoke model not only as a means of cost reduction, but as an enabler of future process improvements. By improving the underlying cost structure in a simpler distribution model organizations can both deliver immediate savings and set themselves up for future process improvements.

References

Agrawal, R., & Srikant, R. (1994). Fast Algorithms for Mining Association Rules. In I. A. Center (Ed.), *Proceedings of the 20th VLDB Conference*, (pp. 487-499). Santiago, Chile.

Alberta Energy System Operator. (2020). *Delivered cost of electricity report.* Calgary, AB: AESO. Retrieved October 14, 2021, from https://www.aeso.ca/assets/Uploads/AESO-Delivered-Cost-of-Electricity-Report-FINAL-31May2020.pdf

Bailey, G. J., & Helms, M. M. (2007). MRO inventory reduction—challenges and management: a case study of the Tennessee Valley Authority, Production Planning and Control. *Production Planning and Control*, 18:3, 261-270,. doi: https://doi.org/10.1080/09537280601127351

Basten, R., & van Houtum, G. (2014). System-oriented inventory models for spare parts. *Surveys in Operations Research and Management Science, 19*(1), 34-55. doi:10.1016/j.sorms.2014.05.002

Bechtel, C., & Patterson, J. (1997). MRO Partnerships: A Case Study. *International Journal of Purchasing and Materials*, 18-23.

Chen, J., Gushikhin, O., Finkenstaedt, W., & Liu, Y.-N. (2019). Maintenance, Repair, and Operations Parts Inventory Management in the Era of Industry 4.0. *International Federation of Automatic Control* (pp. 171-176). Dearborn, MI: IFAC PapersOnLine.

Coupa. (2021). Supply Chain Guru X. *Demand Classification Help Documentation, Version 31.1, 4.0.7950.17907*. Coupa.

García-Benito, J., & Martín-Peña, M.-L. (2020, September 10). A redistribution model with minimum backorders of spare parts: A proposal for the defence sector. *European Journal of Operational Research, 291*(1), 178-193. doi:10.1016/j.ejor.2020.09.014

Ghodrati, B., & Kumar, U. (2005). Reliability and operating environment-based spare parts estimation approach: A case study in Kiruna Mine, Sweden. *Journal of Quality in Maintenance Engineering, 11*(2), 169-184. doi:http://dx.doi.org/10.1108/13552510510601366

Gilbert, J. P., & Finch, B. J. (1985, November 1). Maintenance Management: Keeping Up With Production's Changing Trends and Technologies. *Journal of Operations Management, 6*(1), 1-12. doi:10.1016/0272-6963(85)90030-0

Hauser, C., & Sandoval, E. (2021, July 14). Death Toll From Texas Winter Storm Continues to Rise. *New York Times*.

IEA. (2000-2020). *Global end-use spending on energy.* Paris: IEA. Retrieved October 13, 2021, from https://www.iea.org/data-and-statistics/charts/global-end-use-spending-on-energy-2000-2020

Kennedy, W., Patterson, J., & Fredendall, L. D. (2002). An overview of recent literature on spare parts inventories. *International Journal of Production Economics*, 201-215.

Kukreja, A., Schmidt, C. P., & Miller, D. M. (2001). Stocking Decisions for Low-Usage Items in a Multilocation Inventory System. *Management Science, 47*(10), 1371-1383. Retrieved from http://www.jstor.org/stable/822492

Logistics, D. o., & Support, D. o. (2021, October 1). MIT Capstone - Sponsor Introductions. (J. Maen, & N. Winters, Interviewers)

Logistics, D., & Support, D. (2022, April 11). MIT Capstone - Findings Review. (J. Maen, & N. Winters, Interviewers)

Mann, L. (1966). Toward a Systematic Maintenance Program. *The Journal of Industrial Engineering 17 (9)*, 461-473.

Muckstadt, J. A. (1973). A Model for a Multi-Item, Multi-Echelon, Multi-Indenture Inventory System. *Management Science, 20*(4), 472-481. Retrieved October 30, 2021, from http://www.jstor.org/stable/2629628

Mulvey, J. M., Vanderbei, R. J., & Zenios, S. A. (1995, Mar-Apr). Robust Optimization of Large-Scale Systems. *Operations Research, 43*(2), 264-281. Retrieved October 29, 2021, from http://links.jstor.org/sici?sici=0030-364X%28199503%2F04%2943%3A2%3C264%3AROOLS%3E2.0.CO%3B2-H

S&C. (2018). *S&C's 2018 State of Commercial & Industrial Power Reliability Report.* Chicago: S&C in collaboration with Frost & Sullivan. Retrieved October 11, 2021, from https://www.sandc.com/globalassets/sac-electric/documents/sharepoint/documents---all-documents/technical-paper-100-t120.pdf

Schroder, R. (2004). Minding the Store. *Power Engineering*, 61-63.

Schuh, P., Schneider, D., Funke, L., & Tracht, K. (2015). Cost-optimal spare parts inventory planning for wind energy systems. *Logist. Res.*, 8, 4.

Sherbrooke, C. C. (1966). *METRIC A Multi-Echelon Technique for Recoverable Item Control.* Santa Monica: The RAND Corporation.

Sherbrooke, C. C. (1986, April 1). VARI-METRIC: Improved Approximations for Multi-Indenture, Multi-Echelon Availability Models. *Operations Research, 34*(2), 311-319. doi:https://doi.org/10.1287/opre.34.2.311

Sherbrooke, C. C. (1992). Multiechelon inventory systems with lateral supply. *Naval Research Logistics (NRL), 39*(1), 29-40. doi:10.1002/1520-6750(199202)39:1<29::AID-NAV3220390103>3.0.CO;2-L

Thonemann, U., Brown, A., & Hausman, W. (2002, September 1). Easy Quantification of Improved Spare Parts Inventory Policies. *Management Science*, 1213-1225. doi:10.1287/mnsc.48.9.1213.173

U.S. Bureau of Economic Analysis. (n.d.). Gross Domestic Product. *retrieved from St. Louis Federal Reserve*. Retrieved November 27, 2021, from https://fred.stlouisfed.org/series/GDP

US Energy Information Administration. (2019). Table E8. Primary Energy, Electricity, and Total Energy Expenditure Estimates. *State Energy Data 2019: Prices and Expenditures*, 1. Retrieved 11 28, 2021, from https://www.eia.gov/state/seds/sep_sum/html/pdf/sum_ex_tot.pdf

US Energy Information Administration. (2020). Electricity explained: Factors affecting electricity prices. Retrieved October 13, 2021, from https://www.eia.gov/energyexplained/electricity/prices-and-factors-affecting-prices.php

US Energy Information Administration. (2021, November 3). *Electricity Explained*. Retrieved November 10, 2021, from US Energy Information Administration: https://www.eia.gov/energyexplained/electricity/delivery-to-consumers.php

Weeks, J. (2010, April 28). What is the Smart Grid. *U.S. Electrical Grid Undergoes Massive Transition to Connect to Renewables*. Boston, Massachuessets, USA. Retrieved from SCIENTIFIC AMERICAN, A DIVISION OF SPRINGER NATURE AMERICA, INC.: https://www.scientificamerican.com/article/what-is-the-smart-grid/

Yoder, B. E. (2013). *Developing a Strategic Roadmap for Supply Chain Process Improvement in a Regulated Utility.* Cambridge: Massachusetts Institute of Technology.