Stratification-Based Inventory Optimization to Reduce Network Costs

by

Victoria Wright

and

Waleed Shahid

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

Inventory is a major component of the working capital structure for SLB (formerly Schlumberger), a global technology company that provides energy services to clients. Within this context, optimal positioning of inventory along with the distribution network is critical for smooth operations. As a pilot location, the United States warehouse network was selected to run the project. This project used a model created via a multi-step approach to handle inventory distribution via SKU-segmentation based business rules, combined with mixed integer linear programming for network optimization. The model, once implemented, would result in 1.7 million dollars reduction in inventory requirements and enable the company to move from a disconnected multi-node network to a connected hub-and-spoke model. This will also result in a subsequent reduction in working capital and improve free cash flow. Our recommendation is to implement this model and consider other SLB regions for similar projects.

Capstone Advisor: Dr. Ilya Jackson

Title: Research Scientist

TABLE OF CONTENTS

1 INTRODUCTION

1.1 Background

SLB (formerly Schlumberger) is a global technology company that provides energy services and solutions to clients across more than 100 countries and has employees representing almost twice as many nationalities. It is a global technology leader with offerings that range from traditional upstream oil and gas services to new energy and decarbonization initiatives. In an ever-evolving energy landscape, the company remains focused on optimizing its bottom line to ensure sustainable and profitable growth (Schlumberger Limited, n.d.). Inventory is a major component of SLB's working capital structure. As of December 31, 2023, SLB reported 4.387 billion USD in consolidated inventories, slightly decreasing to 4.375 billion in 2024 (Schlumberger Limited, 2025). Within this context, having an optimized distribution network is critical. Given SLB's global scale, continuous innovation within its supply chain footprint is essential to maintaining operational efficiency, product quality, and safety standards while supporting long-term sustainability.

1.2 Motivation

Product availability remains one of the most persistent challenges within supply chain management. Still, product placement, specifically the geographic positioning of inventory and its accessibility to customers, is just as critical in shaping service levels and total cost. Within SLB, this balance is complex due to the company's large footprint and diverse product portfolio. Currently, SLB applies an undifferentiated approach across a broad range of stock-keeping units (SKUs), with limited consideration for consumption variability or proximity to customer demand. This approach can result in the overutilization of certain warehouses and the underutilization of others. Moreover, when high-demand products are positioned farther from key customers, logistics costs rise, leading to an increased total cost to serve.

Existing research supports the importance of differentiated and data-driven approaches to inventory positioning. Studies on inventory-placement optimization have shown that aligning product location with demand variability and lead times can substantially reduce holding costs (Svoboda & Minner, 2026). Similarly, network design research demonstrates that optimizing both inventory and distribution decisions together leads to measurable reductions in transportation expenses and stock duplication (Allen & Ishida, 2020). These studies collectively highlight that inventory placement should not follow a static model, but should instead evolve dynamically based on service criticality, product valuation, demand variability, and proximity to end users.

This project builds upon the evidence in prior studies and applies these principles within the context of SLB's U.S. Land Geo Unit (an internal SLB geography zone). By incorporating SKU segmentation, velocity analysis, and consumption trends, this project seeks to establish a differentiated inventory-positioning model that improves cost efficiency, optimizes warehouse balance, and reduces emissions within logistics flows while maintaining responsiveness.

1.3 Problem Statement and Key Questions

The primary issue addressed in this project is the inefficient utilization of SLB's warehouse and distribution network. The project will target certain warehouses within the U.S. Land Geo Unit.

At present, SLB applies a uniform inventory strategy across all SKUs, regardless of demand behavior or other relevant metrics. This results in:

1. High cost to serve for slow-moving items.
2. Stock duplication across low-throughput but business-critical warehouses.
3. Underutilized capacity in smaller or remote sites.

There is currently no differentiated approach to inventory positioning based on valuation and consumption/demand variability. Consequently, inefficiencies, unnecessary costs, and scalability constraints persist.

The following key questions guide this project:

1. How can SLB optimize its inventory positioning strategy by redefining SKU segmentation based on consumption variability and geographic proximity to customers to achieve a measurable reduction in cost to serve?
2. Should SLB restructure its physical warehouse network by introducing a central hub to serve as a flow-through point for regional warehouses, and what would be the resulting impact on total inventory on hand, logistics costs, emissions, and overall service levels?

1.4 Project Goals and Expected Outcomes

The goal of this project is to develop a differentiated inventory positioning strategy based on updated SKU segmentation and to evaluate the cost and service impacts of redesigning the warehouse network. The analysis examines the introduction of a centralized hub-and-spoke structure to reduce total inventory on hand while maintaining service levels. By consolidating inventory and optimizing material flows, the project aims to lower cost to serve and improve network utilization. Upon implementation, SLB is expected to achieve a more balanced, cost-efficient, and transparent inventory network across its U.S. Land Geo Unit with clear reduction in inventory requirements and working capital.

2 STATE OF THE PRACTICE

Our main objective is to determine how SLB can optimize its inventory positioning and warehouse distribution network to minimize the total cost to serve its business requirements. The objective function will be driven by two primary inputs: inventory stratification based on consumption trends (both consumption quantity and demand variability) and the current geographical layout of the warehouse network.

To address these challenges, this study investigates three primary areas. First, we examined inventory valuation and distribution network complexities to identify systemic inefficiencies. Second, we evaluated SKU categorization methodologies to determine how to better categorize inventory based on demand patterns. Finally, we researched network optimization techniques to identify strategies for reducing the total cost to serve while maintaining required service levels.

2.1 Current Network Challenges

2.1.1 Categorization and Management of a Large List of Inventory SKUs

In a large inventory system, which covers hundreds to thousands of SKUs, it is a challenge to track each item. Especially when inventory is stored at multiple locations and handled by different systems, this can result in poor tracking and visibility. This lack of transparency prevents real-time inventory valuation and often results in "dead stock" or "stock-outs," as the system cannot dynamically reallocate assets based on shifting demand. For a global sponsor like SLB, this visibility gap directly inflates carrying costs and complicates the goal of achieving a lean distribution model.

2.1.2 Mathematical Complexity of Integrated Network Modeling

The distribution network modeling is commonly addressed using two problems, the facility location problem and the vehicle routing problem. Both can be combined to create a new model that could generate optimal solutions, but the complexity of the said model would require more time and resources to arrive at the optimal network solution. While heuristic solutions exist as an alternative, they typically yield suboptimal results (Arevalo-Ascanio et al., 2024).

2.1.3 Stochastic Demand and Scenario Proliferation

A third critical challenge is demand variability, which introduces uncertainty into what would otherwise be a deterministic model. Incorporating variability requires the creation of stochastic models that account for thousands of potential "what-if" scenarios. Solving these models "as is “meaning, attempting to find a single optimal solution across all possible demand fluctuations is computationally prohibitive because of the curse of dimensionality (Spantidakis, 2022). Without advanced segmentation (such as grouping SKUs by consumption patterns), the network remains unoptimized, forced to either carry expensive excess safety stock or suffer high expedited shipping costs to meet service level agreements.

2.2 Inventory Categorization Techniques

Optimization of a large volume of inventory SKUs requires some key methodologies; one of the main ones used is the concept of stratification/ classification. This concept required grouping of SKUs into predefined categories. This allows for more specific control over each category. It also allows for efficient management of the inventory, since handling a few categories versus a large number of SKUs is achievable with fewer resources (Millstein et al., 2014).

The Pareto Principle serves as the mathematical foundation for this stratification. Often referred to as the "80/20 rule," this principle suggests that a disproportionate share of a system's total value or volume is driven by a small percentage of its components. In the context of inventory management, this means that approximately 80% of total inventory value or consumption is typically derived from only 20% of the SKUs (Pyzdek, 2021). By applying this logic, SLB can prioritize "Class A" items that drive the highest cost or service impact, ensuring that management efforts are focused where they yield the greatest measurable reduction in the cost to serve.

The most common technique in the industry is the ABC classification, which only accounts for a single dimension of the SKU behavior, typically its valuation. Due to this limitation, ABC analysis as a standalone technique is not enough for the efficient planning of inventory, and the results get more exaggerated if other dimensions (for example, demand irregularity) have more impact. Other classification techniques that can be considered include FSN, XYZ, SDE, and HML (Wang et al., 2018) as illustrated in Table 1.

Method	Characteristic	Category
ABC	Annual value $	A (high $$$), B (Medium $$), C (Low $)
FSN	Demand volume	F (Fast), S (Slow), N (Normal)
XYZ	Demand frequency	X (Regular), Y (Fluctuating), Z (Irregular)
SDE	Lead time	S (Scarce), D (Difficult), E (Easy)
HML	Unit value $	H (High $$$), M (Medium $$), L (Low $)

Table 1: Traditional Inventory Classification Approaches (Wang et al., 2018)

In ABC classification, items are classified by either value or volume of consumption. Typically, a few items account for the highest volume/value; a middle category might have a medium level; and the majority of items account for the smallest valuation or volume of movement. These three categories are represented as A, B and C. Another categorization is to apply the ABC analysis on top of the Pareto split, where 20% of the items can be assigned to any one of the A, B or C categories and the remaining items can be assigned to the remaining two categories. Ultimately, the specific thresholds used to define ABC categories are not universal; rather, they vary across the literature depending on the operational context (Millstein et al., 2014).

While ABC categorization applies a certain level of inventory management, it has clear limitations, as it does not account for the demand variability of the inventory items. To help offset this an additional level, XYZ analysis can be applied. A combination of ABC and XYZ is used to optimize inventory positions, resulting in a reduction in cost to serve (Stojanović & Regodić, 2023). The XYZ analysis allows for a secondary categorization of SKUs, which can be based on demand/consumption patterns. The three letters define varying demand trends (Nowotyńska, 2013):

- X is for SKUs with stable demand.
- Y is for demand with medium-level fluctuations.
- Z is for demand that is irregular.

The integration of value and predictability metrics provides a more robust framework for inventory optimization than single-criterion models. A primary method for achieving this is the integrated multicriteria ABC-XYZ method, which allows for the categorization of SKUs into a matrix of nine distinct pairs. This approach enables the development of tailored management strategies for each segment based on clearly defined consumption behaviors (Stojanović & Regodić, 2023).

Under this integrated model, the ABC component segments items by their total value or contribution to revenue, while the XYZ component measures demand stability and forecast accuracy. By merging these two perspectives, a firm can distinguish between high-value items with stable demand (AX) which should be managed with lean, just-in-time replenishment and high-value items with erratic demand (AZ), which require higher safety stock levels to mitigate the risk of stockouts.

To further refine this optimization, the ABC-XYZ matrix can be supplemented by other specialized classification frameworks such as FSN Analysis (Fast, Slow, and Non-moving): This method classifies inventory based on the annual turnover ratio. These categories are applied in different ways to optimize warehouse flow; for instance, "Fast" items are positioned for high accessibility, while "Non-moving" items are flagged for liquidation to recover capital (Putri & Gabriel, 2024).

2.3 Distribution Network Optimization Strategies

Aggregation of existing warehouse networks is a key strategy for tying into the stratification approach discussed above. This aggregation can help an organization reduce different types of expenses like transportation, inventory holding and warehousing fixed costs by leveraging centralization. Other benefits that can be reaped are asset utilization and volume discounts on shipments. However, it's important to consider that absolute consolidation, which refers to the centralized pooling of inventory into a single or minimal number of strategic distribution centers, can lead to increased lead times and the deterioration of inventory availability to customers (Melachrinoudis & Min, 2007). This makes the case for a more nuanced approach where consolidation is applied on a case-by-case basis in tandem with stratification. The aggregation technique can be closely tied to the facility location problem.

These aggregation strategies are fundamentally rooted in the Facility Location Problem (FLP), which generally follows either a continuous or discrete modeling approach. Continuous location modeling assumes a facility can be placed anywhere on a map to find a mathematical "sweet spot." A classic example is the Weber Problem, which identifies a single location that minimizes the weighted sum of distances to all customers. Similarly, the Center of Gravity method calculates a central coordinate to minimize transportation costs based on shipment volumes (Caplice, 2015).

In contrast, discrete location modeling evaluates a specific list of candidate sites to determine the best network setup (Melo et al., 2009). Because these choices dictate long-term capital expenditure (CAPEX) and logistics overhead, they are typically solved using Mixed Integer Linear Programming (MILP). These MILP models use an objective function to minimize total delivery costs while respecting real-world constraints like warehouse capacity and service-level agreements (Caplice, 2015).

Modern optimization must also navigate stochastic(uncertain) demand. To handle this, researchers often use fulfillment rules essentially algorithmic "if-then" logic to distill complex, multi-level distribution problems into a single, manageable layer (Spantidakis, 2022).

3 METHODOLOGY

In this section, we will explain the methodology selected for addressing inventory stratification and how it ties into our network optimization approach. We will start with a justification for our methodology selection and then move toward data collection and processing. Then we will formalize the main steps and milestones for our selected methodology and present some preliminary findings. After completing our literature review and understanding of the core challenges within SLB's distribution network, we have determined that the best approach is twofold, focusing on both inventory stratification and facility location for network optimization.

3.1 Methodology Selection

Inventory Stratification

Prior research establishes that SKU classification is inherently multidimensional; methods focusing solely on value, demand variability, or volume fail to account for intermittency and item-level volatility (Wang et al., 2018; Millstein et al., 2014). Using the ABC analysis, we can ensure that the valuation of inventory is taken into consideration. However, combining this with XYZ allows us to incorporate demand variability, which literature shows becomes more exaggerated when consumption patterns are irregular or cyclical (Nowotyńska, 2013; Stojanović & Regodić, 2023). By utilizing this multifaceted ABC–XYZ matrix, we are better able to identify our inventory characteristics and segment across both value and predictability in a more accurate manner (Stojanović & Regodić, 2023). Since SLB currently does not have a central warehouse in place within its Enterprise Resource Planning (ERP) system, this lack of visibility is critical to address. These stratification outputs help guide the identification of both the most suitable centralized warehouse location and the categories of inventory that should be positioned there.

Network Optimization

After our literature review, we determined that we would focus on utilizing the FLP as the network optimization method. The key drivers for this choice are supported by literature, which identifies the FLP as the standard approach when the objective is to determine which candidate warehouses should remain open and how flows should be routed across the network (Caplice, 2015; Melo et al., 2009; Farahani et al., 2019). This aligns with our need to make a binary decision on whether each warehouse should remain open or closed; verify that demand flows are being met; confirm that capacity limits are respected; and minimize costs through improved routing. The objective of our network optimization is to lower the total cost to serve. This includes the fixed cost of operating warehouses, the variable costs associated with warehouse activities, and transportation costs across the network. The decision variables within the FLP consist of binary variables to determine if a facility should be open or closed; flow variables representing movement from supplier to warehouse and warehouse to end user; and inventory ratio criteria that guide stocking balance based on our ABC and XYZ stratification. The constraints in our optimization include capacity limits for each facility, stocking rules derived from the ABC and XYZ matrix, and satisfaction of all demand requirements across the network. Assumptions might be required to simplify the optimization model to reduce the dimensionality and make the problem manageable.

Scenarios to be tested include the following:

- A case where the warehouse count is minimized to determine the most cost-effective number of facilities.
- An iteration using inventory stratification logic that centralizes slow movers with irregular consumption patterns and keeps critical or high-value items closer to end use to maintain service levels (Melachrinoudis & Min, 2007).

3.2 Data Collection and Preparation

We analyzed multiple datasets to deepen our understanding of the inventory systems and logistics movements within SLB, focusing on the U.S. Land Geo Unit. The datasets have been derived from Corporate Information Models (CIMs) that store all company enterprise data.

Dataset 1: Materials Management Corporate Information Model (CIM)

This dataset, drawn from multiple warehouse tracking systems, provides comprehensive visibility into stock on hand, total stock valuation, unit prices, and consumption trends. The key variables included are warehouse IDs and their corresponding latitude and longitude coordinates; inventory on hand, and encompassing material name, material value, on-hand quantity, unit valuation, and historical consumption trends. This data was used not only to understand the current footprint of our warehouse setup, but also to assess inventory health, providing a starting point for prioritization.

Dataset 2: Shipment Corporate Information Model (CIM)

This dataset, drawn from multiple logistics systems, provides comprehensive visibility into freight activities, costs, and key metrics. Key variables include shipment count (by plant and storage location), shipment volume, transit distance, total cost per shipment, and source/destination plant identifiers.

Dataset 3: Facility Data

This dataset provides the geographic and operational context for each warehouse location in the network. Key variables include latitude and longitude of each facility, warehouse type, and partial fixed warehouse costs. While the dataset was designed to capture both fixed and variable operating costs (including annual rent, annual operating costs, and headcount) complete cost data was only available for a subset of sites due to data limitations encountered during the project. To address the data limitations, we developed a network-wide cost proxy using sites with verified cost data. The proxy was calculated by dividing the total confirmed annual operating cost across those sites by their combined throughput quantity, yielding a single estimated cost per unit handled. This rate was then applied uniformly across all warehouse locations to impute operating costs where direct data was unavailable. While this approach introduces a simplifying assumption (that operating cost per unit is consistent across facility types and geographies) it provided a practical and internally consistent basis for cost estimation given the constraints of the available data. The resulting per-unit rate was multiplied by each facility's 12-month consumption quantity to estimate its annual throughput cost contribution. The limitation of this approach is discussed further in Section 4.3.

3.3 Model Application

Our model consists of multiple key steps, all of which are illustrated in subsections.

3.3.1 Scope

In the project scope, we outline the exact geographical zone as a U.S. Land Geo Unit for SLB and specifically the Standard Warehouses. Given the volume of data across SLB's full warehouse network, the scope was intentionally limited to Standard Warehouses as the first phase of analysis. This decision was based on the assumption that inventory behavior patterns within Standard Warehouses are sufficiently consistent to model independently, and that the methodology can be extended as a filter to other warehouse types in future iterations since it is built to be scalable.

3.3.2 Data Working

The data preparation phase required significant effort to reconcile the structure of SLB's internal data models with the requirements of the analysis. A key challenge was that SLB operates multiple Stock Locations and Centers (SLOCs) that can be tied to a single physical warehouse location. This created a many-to-many relationship in the raw data that would have distorted both the cost calculations and the optimization model. To resolve this, the data was restructured into a many-to-one relationship, consolidating multiple SLOCs to their corresponding warehouse location. This simplification was necessary not only for analytical clarity but as a prerequisite for the MILP model to run successfully. Beyond this structural fix, the dataset required standard cleaning procedures, including the removal of outliers, correction of data entry errors, and alignment of column formats across the four source datasets before a unified data model could be constructed.

3.3.3 Inventory Positioning

We conducted Inventory stratification in two sequential layers. First, we applied ABC classification to the full SKU list to rank items by their total stock consumption value over the last 12 months, identifying which items represented the most significant financial exposure to the network. Then we layered XYZ classification on top of the ABC output to characterize demand variability for each item, distinguishing between items with stable, predictable consumption patterns and those with erratic or intermittent demand. The combination of these two dimensions produced an ABC-XYZ matrix that served as the foundation for assigning inventory policies across the network. The strategic intent of this layering was to identify which SKUs were strong candidates for centralization. The underlying logic was that slow-moving, low-variability items carry less service risk when consolidated into a central hub since there is sufficient lead time to plan replenishment without impacting field operations. In contrast, high-value, high-demand items with unpredictable consumption were flagged for local stocking, as delayed availability could directly affect operational continuity. The resulting inventory positioning policies were developed in Table 2.

Plant Level Classification Logic				
Rule	**Plant Class**	**Examples**	**Positioning Decision**	**Inventory Buffer**
Rule 1	Ends in X (Any)	AX, BX, CX	Local Stock – Full SS	Full Independent Buffer at Plant
Rule 2	Ends in Y (Any)	AY, BY, CY	Local Stock – Reduced SS	Lean Local Buffer (Tier 2 SLA)
Rule 3	Ends in Z (Any)	AZ, BZ, CZ	Central Stock Only	Centralized Pooled Buffer
Rule 4	—	Zero Consumption	No Stock / Review	Flagged for Dead Stock Review
X — Stable, High-Velocity Demand (Kept Local) | Y — Variable Demand (Kept Local, Reduced SS) | Z — Slow-Moving or Intermittent Demand (Centralized)				
$72.4M Identified for Centralization				

Table 2. Inventory Positioning Policies

3.3.4 Network

The baseline network in scope for U.S. Land Geo Unit Standard Warehouses consisted of 52 warehouses with a total stock value of $155 million. Following the application of the inventory positioning policies described above, $72.4 million was identified as a candidate for centralization, forming the basis for network optimization. The optimization model was designed to assign each plant to a regional hub and minimize the total cost of the network. The model runs iteratively to determine the optimal number of hub candidates rather than requiring a fixed hub count as an input. The high-level objective function is defined as:

Total Cost = Freight Cost + Inventory Holding Cost + Hub Upgrade Cost + Service Level Penalty

The detailed objective function along with functional constraints can be illustrated mathematically as follows:

$$\text{Minimize Total Cost} = \sum_{i \in I} \sum_{j \in J} \alpha W_i d_{ij} x_{ij} + \sum_{j \in J} \pi(K) y_j + \sum_{j \in J} F\, z_j + \sum_{i \in I} P\, s_i$$

The first term in the detailed objective function represents freight cost, the second term represents inventory holding cost derived from the Square-Root Law of inventory pooling, the third term represents fixed hub upgrade costs, and the final term penalizes service-level violations.

The inventory holding cost per hub is derived from the Square-Root Law of inventory pooling, which states that safety stock scales with √(k/n) as n nodes consolidate into k hubs. To account for demand correlation across SLB's network driven by shared oil-price cycles and drilling seasonality a correlation-adjusted factor (ρ) is applied:

$$\pi(K) = \frac{C_0\left[\sqrt{K/n_0} \cdot (1 - \rho) + \rho\right]}{K}, \quad K = \sum_{j \in J} y_j$$

The function $\pi(K)$represents the inventory holding cost per hub as a function of the number of hubs selected in the network. It captures the well-established Square-Root Law of inventory pooling, which states that total safety stock in a distributed system increases with the square root of the number of stocking locations.

The constraints for this model are listed as follows. These constraints ensure that each plant is assigned to exactly one open hub, that hub activation and regional designation decisions are logically consistent, and that service-level distance requirements are respected with controlled violations captured through penalty variables. Table 3 represents the descriptions of all the variables used in the model.

$$\sum_{j \in J} x_{ij} = 1 \quad \forall i \in I$$

$$x_{ij} \le y_j \quad \forall i \in I, j \in J \qquad z_j \le y_j \quad \forall j \in J$$

$$\sum_{i \in I, i \neq j} x_{ij} \le |I| z_j \quad \forall j \in J$$

$$\sum_{j \in J} d_{ij} x_{ij} \le SLA_i + s_i \quad \forall i \in I$$

$$x_{ij}, y_j, z_j \in \{0,1\}, \qquad s_i \ge 0$$

Symbol	Type	Description	Value / Notes
I	Set	Plants / warehouses	—
J	Set	Candidate hubs	—
x_{ij}	Binary	Equals 1 if plant i is assigned to hub j	—
y_j	Binary	Equals 1 if hub j is opened	—
z_j	Binary	Equals 1 if hub j is regional	—
s_i	Continuous ≥ 0	Service Level Agreement (SLA) distance violation at plant i	—
d_{ij}	Continuous	Distance from plant i to hub j	miles
W_i	Continuous	Annual consumption value at plant i	USD
α	Constant	Freight cost coefficient	0.000075
F	Constant	Hub upgrade cost	$50,000
P	Constant	SLA penalty coefficient	Large penalty
SLA_i	Constant	Maximum service distance for plant i	Tier-based
C_0	Constant	Baseline inventory holding cost	Derived
n_0	Constant	Original number of stocking locations	Fixed
K	Integer	Number of open hubs	$\sum_{j\in J} y_j$
$\pi(K)$	Function	Inventory cost per hub	Square-Root Law
ρ	Constant	Cross-location demand correlation coefficient	0.3

Table 3. Model Variables and Parameters

Key assumptions and parameters used in the model are defined as follows. A transportation cost coefficient $\alpha = 0.000075$ was applied as a proxy for freight cost per unit value-mile, enabling transportation cost to be modeled as a value-weighted distance function. An annual inventory holding rate of $H = 20\%$ was used to estimate the cost of carrying inventory across the network. A fixed hub upgrade cost of $F = \$50{,}000$ was assigned to any plant designated as a regional hub, reflecting the assumption that operational and infrastructure changes would be required at those locations in the absence of detailed capacity data. Service tiers were assigned at the plant level based on the inventory positioning policies established during stratification (Table 4), with each plant-specific service level parameter SLA_i defining the maximum allowable distance inventory may travel from a hub to a served plant, as illustrated in Figure 1. To enforce compliance with these service-level distance limits, a large penalty coefficient P was applied to any violation of the SLA constraint, ensuring that service breaches occur only when no feasible alternative exists. The baseline inventory holding cost C_0 was derived by applying the holding rate H to the total inventory value eligible for centralization under the baseline network structure. The parameter n_0 represents the original number of stocking locations in the network and serves as the reference point for the Square-Root Law of inventory pooling, which governs how total inventory holding cost scales as the number of hubs K changes in the optimized network. A demand correlation parameter ρ adjusts the pooling factor to $\sqrt{K/n_0} \cdot (1-\rho) + \rho$, accounting for cross-location demand co-movement across SLB's network. A base-case value of ρ = 0.30 is applied.

Service Tier Classification & Radius Policy				
Service Tier	**Classification Logic**	**Max Radius**	**Delivery Strategy**	**SLA Penalty**
Tier 1: Rapid	≥50% Inventory in Local Stock – Full SS	150 Miles	High-Velocity Local Fulfillment	Extremely High — Strictly Enforced
Tier 2: Standard	≥50% Inventory in Local Stock – Reduced SS	300 Miles	Balanced Regional Coverage	High — Enforced
Tier 3: Economy	All Other Active Stocking Profiles	600 Miles	Consolidated Hub Fulfillment	Moderate
Tier 4: Passive	Classified by Total Stock Value — Low Velocity	1,500 Miles	Long-Haul / On-Demand	Low — Flexible
Radius Limits Enforced via High SLA Penalty in Objective Function — Inventory Relocation Stays within Operationally Appropriate Bounds per Tier				

Table 4. Service Tiers at the Plant Level

The model applies a high service level penalty to enforce compliance with tier-based radius constraints, ensuring that inventory is not served from distances that are operationally inappropriate for the material type. For example, Tier 1 high-velocity, high-value items are prevented from being served from distant hubs, with the penalty structure enforcing this boundary mathematically. Figure 1 illustrates the service level radius and hub selection logic. In this example, warehouse 2016 was assigned a 600-mile service radius based on its inventory positioning tier, and the model selected the nearest qualifying hub while respecting this constraint. Network optimization defines feasible network paths and stocking policies rather than physically relocating inventory.

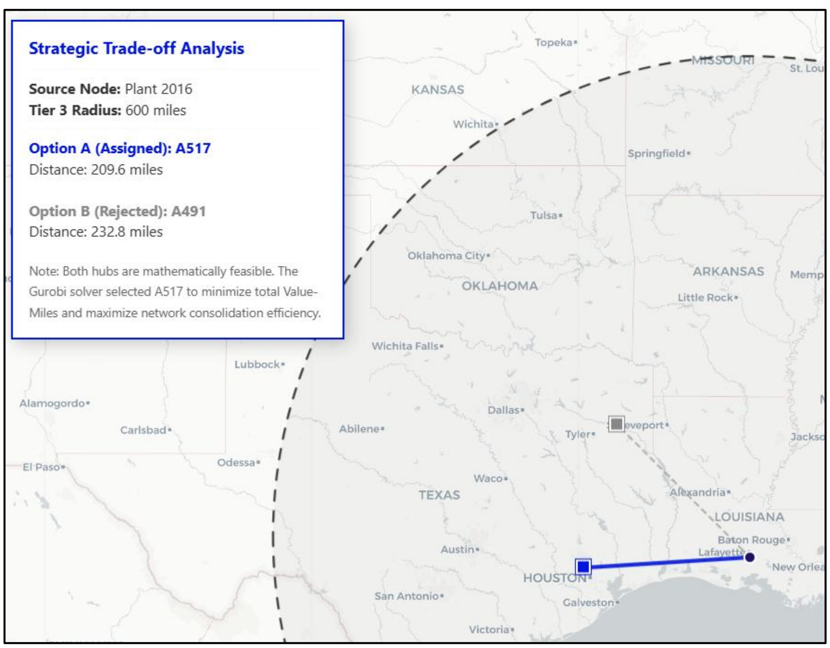

Figure 1: Illustration of Radius Rule in Network

4 RESULTS AND DISCUSSION

In this section, we will present and discuss the results of the model. We will start by presenting some initial statistics and visualizations on the stratification approach, followed by a before and after of the network optimization model and conclude with a summary of the new proposed network and savings potential.

4.1 Summary Statistics

Table 5 provides a summary of our initial data.

Metric	UOM	Quantity
Total Warehouses	EA	52
Total SKUs (Material Numbers)	EA	671
Total Stock Value	USD	155,388,282
Total Last 12M Consumption	USD	371,345,930

Table 5: Summary Statistics of Geography

It illustrates the total footprint we have selected for this project, namely all Standard Warehouses in U.S. Land Geo Unit for SLB. It is important to note there are other types of warehouses and locations (such as manufacturing centers) which exist in this region, but they were not considered as part of the same network, because they are governed by their own policies. After evaluating our warehouse footprint, we took a closer look at SKUs that haven't moved in the last year. Identifying these 'No Consumption' parts is vital for our policy alignment. Following the methodology, we established earlier, we are focusing our optimization and centralization strategies strictly on parts with active demand. The logic is that parts with no consumption would not yield any safety stock reduction after pooling. You can see the breakdown of these excluded items in Table 6.

Inventory Status	Quantity of SKU	Stock Value (USD)	% Value
Parts With Consumption	460	114,426,449	73.64%
Parts With No Consumption	211	40,961,833	26.36%
Total Global Inventory	671	155,388,282	100.0%

Table 6: Consumption Summary of SKUs

To maintain a lean and high-functioning central hub, we define non-consumption as any inventory item with zero utilization over a rolling 12-month period. While the long-term goal is total centralization, these stagnant items require a nuanced, independent review process. We must differentiate between truly obsolete stock (which threatens to clutter the hub) and critical-spare components (which may have low turnover but remain vital for plant uptime). Consequently, disposition decisions for these items will be made through collaborative reviews with individual plant leadership to ensure operational readiness isn't sacrificed for storage metrics.

4.2 Stratification and Positioning Application

As illustrated in the methodology section's inventory stratification approach, all SKUs get mapped into ABC-XYZ matrix with respect to the last 12 months of consumption value and coefficient of variation. It is important to note that both these metrics are at plant level, such as the example shown in Figure 2.

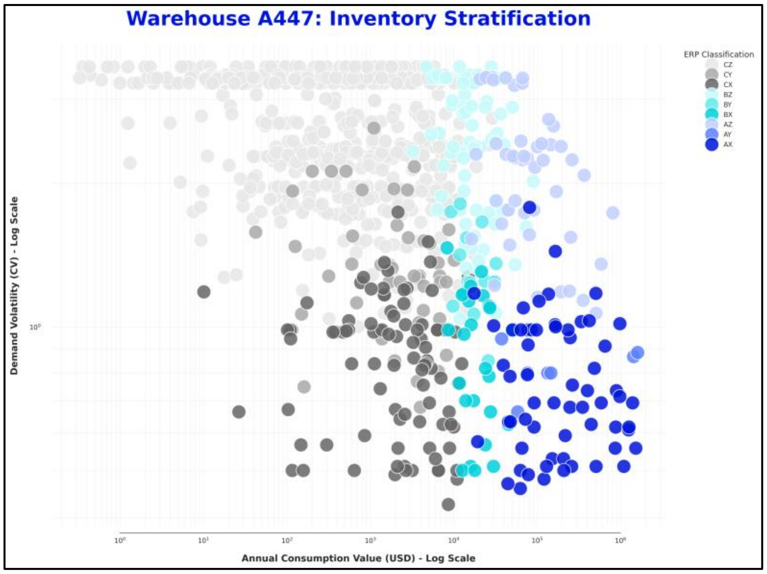

Figure 2: Plant A447 Stratification Visualization

In this example, we do observe that the splits are not showing perfectly due to a delay in policy implementation in the ERP system relative to the trends reflected in the latest consumption history. After application of stratification, we apply the inventory positioning logic as illustrated in the methodology section. This resulted in distribution of stock as per Table 7.

Positioning	Total Stock Value USD	Stock Value % of Total
Central Stock Only	56,046,370	36.07%
Local Stock – Full SS	41,983,810	27.02%
Local Stock – Reduced SS	16,396,269	10.55%
No Stock / Review	40,961,833	26.36%
Grand Total	**155,388,282**	**100.00%**

Table 7: Inventory Distribution by Positioning

We can clearly see that our positioning policy is highlighting around 46% of inventory, which is considered for centralization efforts under categories “Central Stock Only” and “Local Stock – Reduced SS.” We also see that around 26% is the portion with no consumption history, which we also observed in section 4.1 under summary statistics. It is important to note that, of the 46% of inventory identified as suitable for centralization, only 5.7% can be truly pooled, as these SKUs are common across warehouses. Therefore, the savings calculations for this project are based on this revised proportion.

4.3 Network Optimization

For network optimization we will start by sharing Figure 3, which displays the original warehousing network without the application of any network model.

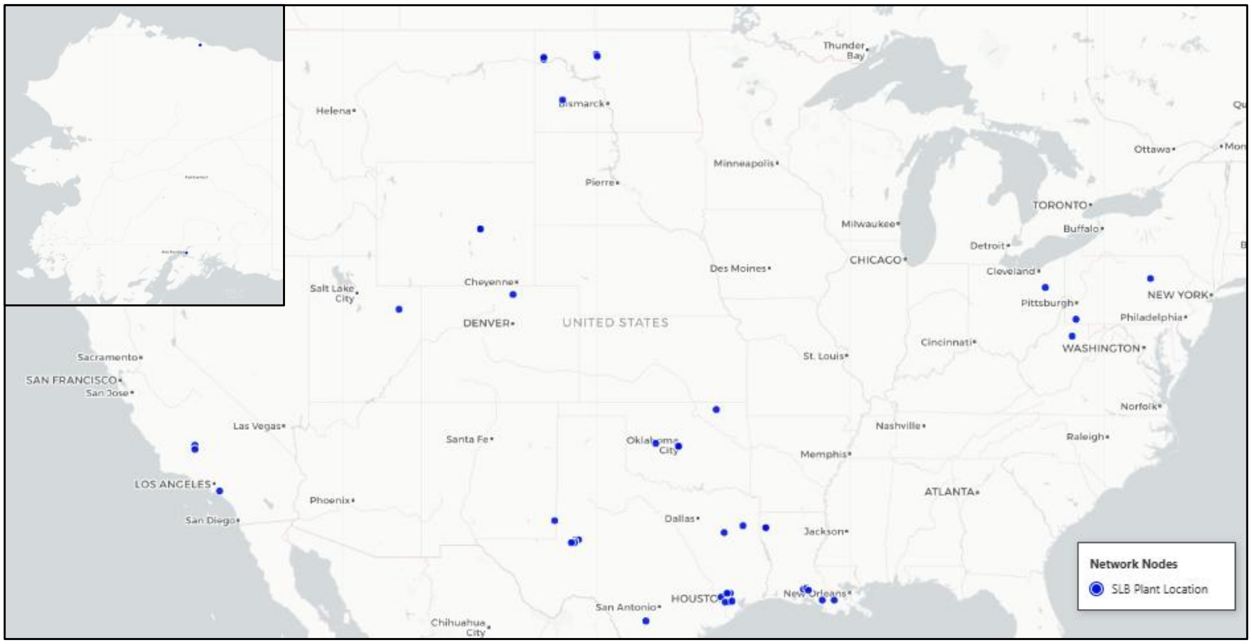

Figure 3: Original Warehouses' Distribution

After application of the stratification matrix along with the positioning strategy and subsequent network optimization, Figure 4 represents the proposed network of hubs as well as the satellite locations they would be serving. As you can see, the model proposes a hub-and-spoke structure, with nine locations upgraded to hub status and the remaining locations being served either from hubs or themselves depending on the inventory positioning strategy. This allows for minimal disruption of existing structure.

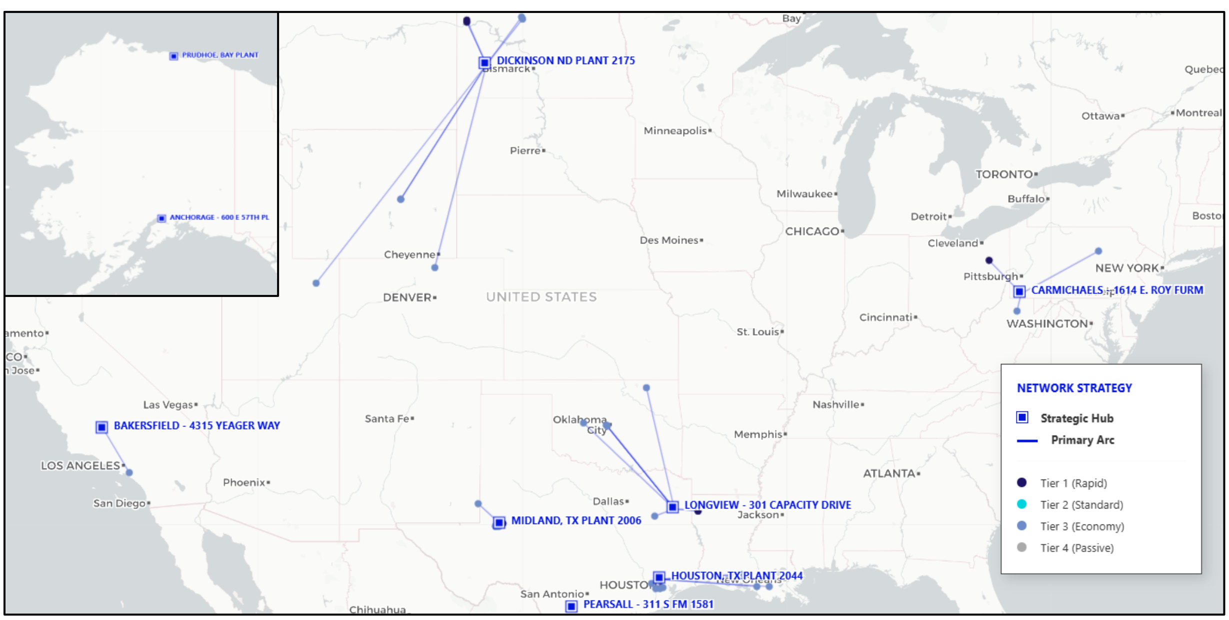

Figure 4: Proposed Network Map

A more zoomed representation of a specific cluster is provided in Figure 5, because Figure 4 does not show all the details.

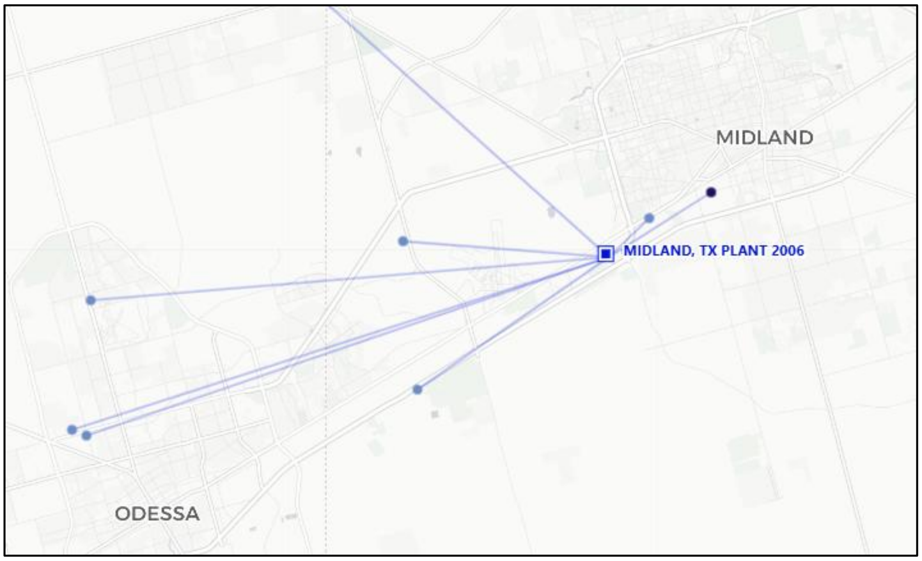

Figure 5: Midland, TX. Proposed Hub Network

4.4 Strategic Summary

A strategic summary of the network results is listed in Table 8. Here, OPEX refers to the operating cost of the total network, calculated using a fixed cost per unit of throughput. This cost is assumed to be the same across all warehouses to keep the analysis streamlined and consistent. The redistribution freight cost is part of the Total Network OPEX cost.

Strategic Metric	Baseline	Optimized	Delta (%)
Total Network Nodes	52	52	-
Active Strategic Hubs	0	9	-
Field Satellites	0	43	-
Full Network Throughput (Qty)	67,091,561	67,091,561	0.0%
Total Inventory ($)	**$155,388,282**	**$153,707,138**	**-1.1%**
- Inv at Hubs ($)	$0	$90,913,151	-
- Inv at Satellites ($)	$0	$62,793,987	-
- Redistribution Freight ($)	$0	$33,903	-

Table 8: Comparative Impact of Hub-and-Spoke Integration on Network Capital and Efficiency

The results clearly demonstrate that pooling inventory across hubs generates substantial capital savings. While the model introduces approximately $33,903 in additional redistribution freight costs due to increased hub-to-satellite movements, these incremental logistics expenses are relatively minor when compared to the financial benefits achieved. Specifically, the centralization strategy enables a reduction in inventory levels of approximately $1.7 million (1.1%), which in turn significantly lowers associated holding costs. This reduction represents the most impactful driver of value in the analysis, far outweighing the additional transportation costs. Overall, the findings confirm that inventory pooling not only improves operational efficiency but also delivers considerable net economic benefits by optimizing capital allocation within the supply chain.

4.5 Limitations

Several limitations should be considered when interpreting the results. Fixed and variable operating cost data were available only for a subset of warehouses, requiring the use of a network-wide cost proxy that applies a uniform cost-per-unit assumption across all sites. Customer destination data were not available, so the model optimizes warehouse-to-warehouse assignments rather than true demand-point relationships, relying on distance-based service radius rules. In addition, service-level radius limits were based on internal operational policies rather than contractual SLAs, and freight cost parameters were benchmarked from industry sources rather than derived from actual SLB shipment records.

5 RECOMMENDATIONS

This section outlines the key recommendations that emerge from our analysis, focusing on the practical steps SLB can take to capture the benefits identified in the study. While the optimization model provides a strong foundation for reshaping the network, successful implementation will require both strategic decisions at the management level and targeted follow-up work to refine assumptions, validate inputs, and ensure operational readiness. The recommendations that follow highlight the most important actions to support a smooth transition toward a more efficient, centralized, and cost-effective warehouse network.

5.1 Management Recommendations

To fully realize the benefits highlighted in this project, SLB’s leadership should take a structured and phased approach to implementation. The recommendations below outline the key measures required to support a smooth transition. In addition, they are aligned with answers to the two questions postulated in section 1.3 of the paper.

- Establish clear ownership of the new network design. Assign responsibility across supply chain, operations, and planning teams to ensure consistent adoption of the new inventory-positioning framework.
- Refine and validate cost assumptions with Finance. Update the operating-cost proxy, hub-upgrade assumptions, and freight-rate inputs to ensure financial accuracy before full deployment.
- Formalize the updated SKU segmentation and stocking policies. Integrate the ABC–XYZ segmentation and corresponding positioning logic into standard operating procedures so that they become part of ongoing planning cycles instead of a one-time study.
- Develop a phased implementation roadmap. Sequence the rollout of the hub-and-spoke network, beginning with validation of hub readiness, site capacity, and transition requirements for each affected warehouse.
- Coordinate change management with field operations. Engage field teams early to minimize service disruptions and ensure operational leaders understand how inventory flows will be changing.

- Implement governance for continuous monitoring. Establish regular review cycles to monitor service levels, freight cost trends, inventory performance, and hub utilization, enabling the network to adjust as business conditions evolve.
- Support cross-functional alignment with the domestic logistics function. This feedback will be essential for defining local logistics costs. These costs are projected to rise under a centralized model, as repositioning SKUs will require additional physical inventory movements.

5.2 Future Work

5.2.1 Review of Non-Consumption Inventory

As established in the methodology, no-consumption inventory SKUs with zero usage over the trailing 12 months has been excluded from the centralization model to avoid transferring potential obsolescence to the regional hubs. However, these items require a deliberate plant-level review to distinguish between genuine obsolete stock and critical operational spares. Because this distinction requires local operational judgment rather than data-driven modeling, each plant should coordinate an independent audit with maintenance stakeholders. This ensures that only active, high-value inventory is consolidated, while stagnant stock is either written off or managed locally as emergency spare part inventory.

5.2.2 Updated Facility Cost Data

A second recommendation addresses the cost-proxy limitation used in the model. Operating costs were estimated using a network-wide throughput proxy derived from a subset of facilities, which introduces uncertainty into hub selection and savings estimates. With complete site-level fixed and variable cost data, the model could be re-run to more accurately assess consolidation benefits, particularly for smaller or lower-throughput sites where break-even outcomes are sensitive to cost assumptions. Additionally, a uniform hub upgrade cost of $50,000 was applied due to limited facility condition data; future model iterations should incorporate site-specific capacity and capital requirements to better reflect actual upgrade costs and refine hub selection decisions.

REFERENCES

Allen, M., & Ishida, R. (2020). *B2B omnichannel network design and inventory positioning* [Master's thesis, Massachusetts Institute of Technology]. https://ctl-archive.mit.edu/sites/ctl.mit.edu/files/theses/B2B%20Omnichannel%20Network%20Design%20and%20Inventory%20Positioning.pdf

Arevalo-Ascanio, R., De Meyer, A., Gevaers, R., Guisson, R., & Dewulf, W. (2024). Location-routing problem for integrated supply chain network design with first and last mile: A critical literature review. *Operations and Supply Chain Management: An International Journal, 17*(2), 206–219. https://doi.org/10.31387/OSCM0570423

Caplice, C. (2015). *CTL.SC2x: Supply chain & logistics design – intro to network design.* edX.

Farahani, R. Z., Fallah, S., Ruiz, R., Hosseini, S., & Asgari, N. (2019). OR models in urban service facility location: A critical review of applications and future developments. *European Journal of Operational Research, 276*(1), 1–27. https://doi.org/10.1016/j.ejor.2018.07.036

Melachrinoudis, E., & Min, H. (2007). Redesigning a warehouse network. *European Journal of Operational Research, 176*(1), 210–229. https://doi.org/10.1016/j.ejor.2005.04.034

Melo, M. T., Nickel, S., & Saldanha-da-Gama, F. (2009). Facility location and supply chain management – a review. *European Journal of Operational Research, 196*(2), 401–412. https://doi.org/10.1016/j.ejor.2008.05.007

Millstein, M. A., Yang, L., & Li, H. (2014). Optimizing ABC inventory grouping decisions. *International Journal of Production Economics, 148*, 71–80. https://doi.org/10.1016/j.ijpe.2013.11.007

Nowotyńska, I. (2013). An application of XYZ analysis in company stock management. *Modern Management Review, 18*, 67–76. https://doi.org/10.7862/rz.2013.mmr.7

Putri, I., & Gabriel, D. S. (2024, June). ABC–FSN analysis for inventory management policy: A case study in retail industry. *Proceedings of the 9th North American Conference on Industrial Engineering and Operations Management*. https://doi.org/10.46254/NA09.20240185

Pyzdek, T. (2021). Pareto analysis. In *The lean healthcare handbook*. Springer. https://doi.org/10.1007/978-3-030-69901-7_14

Schlumberger Limited. (n.d.). *Who we are*. https://www.slb.com/about/who-we-are

Schlumberger Limited. (2025). *Form 10-K for the fiscal year ended December 31, 2024*. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/87347/000095017025007638/slb-20241231.htm

Spantidakis, I. (2022). *Constrained inventory optimization on complex warehouse networks* [Doctoral dissertation, Massachusetts Institute of Technology]. MIT DSpace. https://hdl.handle.net/1721.1/147328

Stojanović, M., & Regodić, D. (2023). The significance of the integrated multicriteria ABC–XYZ method for the inventory management process. *Acta Universitatis Oeconomicae, 3*, 1–15. https://acta.uni-obuda.hu/Stojanovic_Regodic_76.pdf

Svoboda, J., & Minner, S. (2026). A data-driven approach for strategic inventory placement in multi-echelon supply networks. *European Journal of Operational Research, 328*(2), 446–459. https://doi.org/10.1016/j.ejor.2025.06.022

Wang, F., Ng, H. Y., & Ng, T. E. (2018, December 16–19). Novel SKU classification approach for autonomous inventory planning. Proceedings of the *IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, 1441–1445. https://doi.org/10.1109/IEEM.2018.8607736