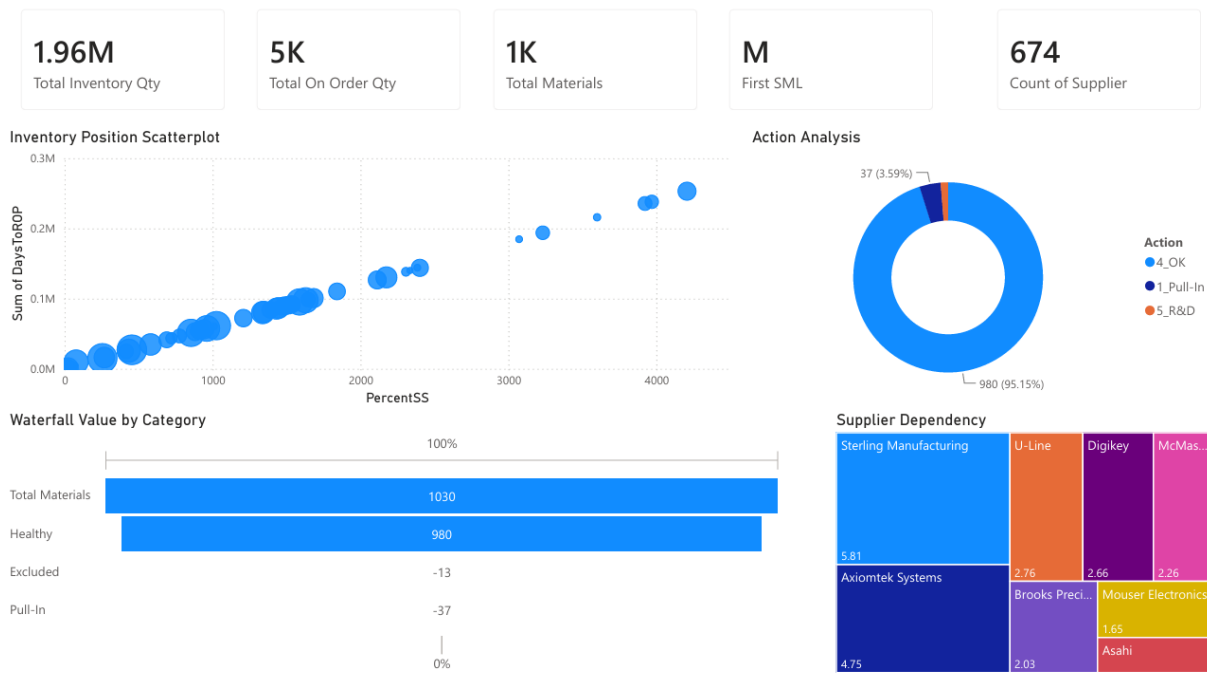


Engineering Operational Decision Systems: Rebuilding a Procurement Planning Engine

By Fatima Aguilar

A case study in systems engineering, inventory optimization, and enterprise process redesign.



An automated inventory planning and decision-support system that integrates SAP data, PowerShell automation, Excel calculation logic, and Power BI dashboards into a single daily decision ritual.

Procurement Heatmap Engine

**AUTOMATING INVENTORY DECISIONS.
EMPOWERING STRATEGIC BUYING.**

OVERVIEW

The Procurement Heatmap Engine is an end-to-end automation solution that transforms how buyers monitor inventory health and make purchasing decisions.

It replaces a manual, error-prone Excel-based process with an automated pipeline that integrates SAP data, cleanses and enriches it through code, applies business rules, and delivers clear, actionable recommendations through a Power BI dashboard and scheduled outputs.

The result is faster insights, consistent execution, and better inventory decisions across the organization.

IMPACT AT A GLANCE



80%+

Reduction in manual work



100%

Standardized logic and outputs



Faster Decisions

Timely, accurate insights for buyers



Data Integrity

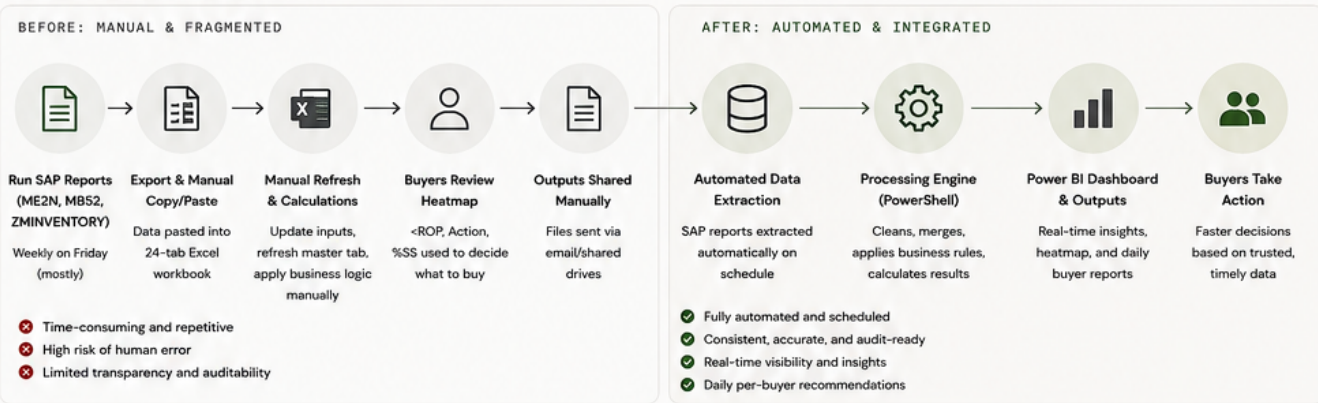
Clean, reliable, and audit-ready data



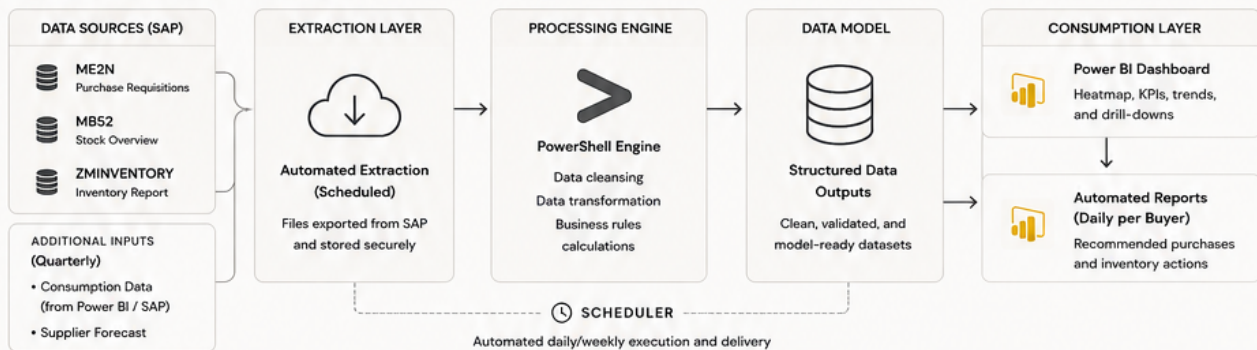
Strategic Impact

Better inventory health, lower risk, and improved service levels

PROCESS: BEFORE VS AFTER



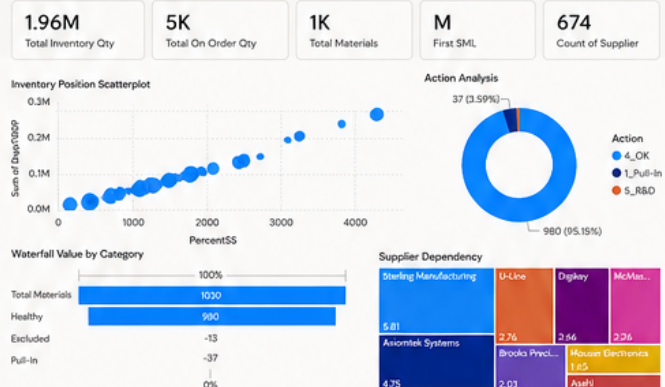
ARCHITECTURE: END-TO-END DATAFLOW



BUSINESS LOGIC (HEATMAP RULES)

<ROP (Inventory Position)	ACTION (Recommended)	KEY CONCEPTS
1 1_<ROP Below Reorder Point	1 1_Pull-in Expedite receipts / bring forward	On Hand Available inventory
2 2_OK Healthy	2 2_<ROP Create PO / Replenish	Reorder Point (ROP) Minimum level to trigger buy
3 3_Exclude R&D / Not Applicable	3 3_Push-Out Delay / Push to later date	Under Inspection Stock in quality hold
4 4_OK No action needed	4 4_OK No action needed	Min / Max Stock Inventory boundaries
5 5_R&D Research & Development	5 5_R&D Research & Development	Safety Stock (SS) Buffer for variability
6 6_Obsolete Obsolete / Do not buy	6 6_Obsolete Obsolete / Do not buy	

DASHBOARD PREVIEW



KEY METRICS



TECHNOLOGY STACK



Introduction

Modern supply chains generate enormous quantities of operational data. Enterprise Resource Planning (ERP) systems continuously record inventory balances, purchase orders, supplier lead times, consumption history, material classifications, and hundreds of additional planning variables. Despite this abundance of information, many operational purchasing decisions continue to rely on fragmented spreadsheets, manual calculations, and human interpretation rather than integrated decision support systems.

Over time, these spreadsheet based planning tools often evolve into highly complex operational platforms. As new requirements emerge, additional worksheets, formulas, helper columns, lookup tables, and manual processes are layered onto existing workbooks. Although these systems may continue producing acceptable business results, they become increasingly difficult to understand, maintain, validate, and extend. The knowledge required to operate them frequently resides with only a small number of experienced users, creating significant operational risk and limiting scalability.

Inventory planning presents a particularly challenging decision environment because every purchasing recommendation depends on the interaction of multiple dynamic variables. Current inventory, future demand, open purchase orders, safety stock requirements, supplier lead times, inventory classifications, cost information, and business exceptions must all be evaluated simultaneously before determining whether material should be expedited, replenished, delayed, excluded, or allowed to remain unchanged. As product portfolios expand into thousands of active materials, manually evaluating these relationships becomes increasingly time consuming and inconsistent.

The objective of this project was not simply to automate an existing spreadsheet. Instead, the goal was to redesign an operational decision system from first principles.

Rather than treating Excel as the computational engine, the project separated business logic from presentation, transforming years of accumulated spreadsheet calculations into a centralized software architecture capable of producing deterministic, auditable, and maintainable outputs. The resulting system, referred to as the Procurement Heatmap Engine, integrates SAP operational data, inventory planning logic, and procurement business rules into an automated pipeline that generates standardized inventory health assessments, purchasing recommendations, Power BI dashboards, and scheduled buyer reports.

The redesign process began with a complete architectural audit of the existing planning workbook. Every worksheet was analyzed according to its functional purpose, dependencies, manual intervention requirements, and contribution to the overall decision making process. This analysis revealed redundant worksheets, obsolete calculations, duplicated business logic, and tightly coupled formula dependencies that had accumulated over several years of incremental development. Instead of reproducing this complexity in software, the project reconstructed the planning process around a modular architecture centered on a single calculation engine.

The final solution replaces distributed Excel formulas, Power Query transformations, and manual worksheet refreshes with a PowerShell based processing engine that performs data aggregation, inventory calculations, business-rule evaluation, and recommendation generation through a centralized and version-controlled codebase. The outputs feed interactive Power BI dashboards and automated reporting workflows that enable procurement professionals to make faster, more consistent, and more transparent inventory decisions.

Beyond improving operational efficiency, this project demonstrates how systems engineering principles can be applied to modernize legacy operational tools. By viewing spreadsheet infrastructure as an engineering system rather than a collection of files, organizations can reduce technical debt, improve decision quality, increase maintainability, and establish a scalable foundation for future analytics, optimization, and artificial intelligence applications.

Understanding the Legacy System

Before any automation could be developed, the existing procurement planning process had to be understood as a complete operational system rather than simply an Excel workbook. Although the planning process appeared to revolve around a single spreadsheet, the workbook had evolved over several years into the organization's primary decision support platform for inventory management. It served as the central location where SAP data, inventory policies, planning assumptions, supplier information, and buyer expertise converged to produce purchasing decisions.

Each week, procurement planners exported multiple reports from SAP, including purchase requisitions (ME2N), inventory balances (MB52), and inventory status reports (ZMINVENTORY). Additional planning inputs, including quarterly consumption data, supplier forecasts, and manually maintained planning parameters, were collected from various business systems and incorporated into the workbook before any calculations could begin. The workbook was then manually refreshed through a sequence of Power Query transformations, Excel formulas, lookup tables, and helper worksheets before buyers could evaluate inventory conditions and determine the appropriate purchasing actions.

Although the process successfully supported procurement operations, it depended heavily on manual execution. Individual planners were responsible for downloading reports, updating worksheets in the correct sequence, validating intermediate calculations, refreshing queries, and interpreting the resulting outputs. Because much of the business logic was embedded directly within Excel formulas, understanding how a recommendation had been produced often required tracing calculations across numerous worksheets.

Rather than immediately automating the workbook, the first phase of the project focused on understanding its architecture.

A complete architectural audit was performed in which every worksheet was cataloged according to its operational purpose, data ownership, dependency relationships, level of manual intervention, automation potential, and contribution to the overall planning process. This analysis transformed what initially appeared to be a collection of spreadsheets into a documented information system whose components could be evaluated individually.

The audit identified approximately twenty-four worksheets that collectively supported the inventory planning process. Each worksheet was classified as either an input source, calculation layer, configuration table, output, or legacy

artifact. The review revealed that while several worksheets remained essential to procurement planning, others had become redundant, obsolete, or duplicated functionality already available elsewhere in the workbook. Some worksheets existed solely to support intermediate calculations, while others contained manually maintained lookup tables or historical logic that was no longer actively used in decision making.

More importantly, the audit exposed a highly coupled architecture. Data flowed through multiple manual transfers before reaching the primary inventory management worksheet. Business rules were distributed across helper columns, Power Query transformations, XLOOKUP functions, nested Excel formulas, and manually maintained logic tables. Because calculations were dispersed throughout the workbook, even relatively small modifications required understanding dependencies spanning numerous worksheets. This complexity increased maintenance effort, reduced transparency, and introduced operational risk whenever business rules or planning assumptions changed.

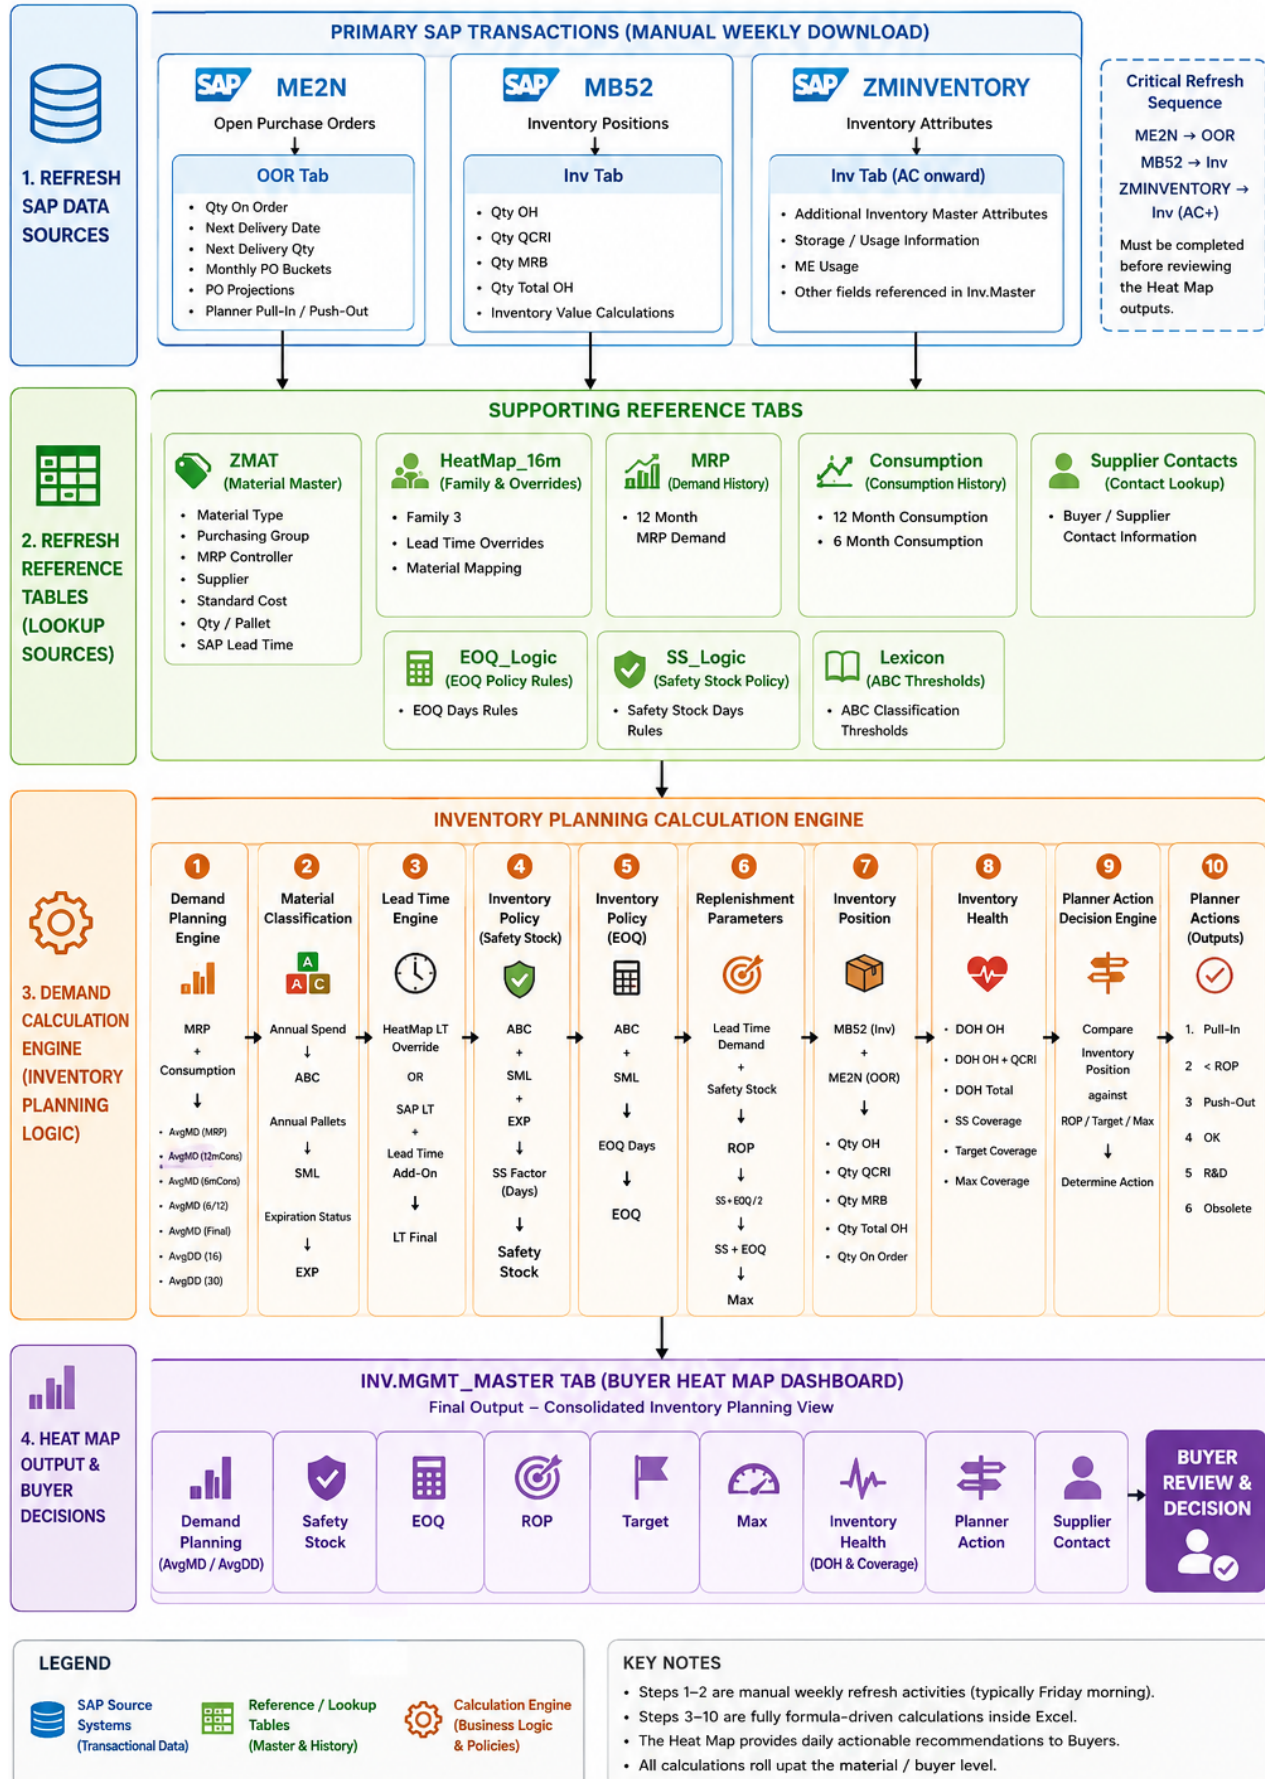
The workbook architecture also reflected a common characteristic of mature operational systems: years of incremental improvements had optimized individual tasks without reconsidering the overall system design. New worksheets had been added as business requirements evolved, formulas had accumulated to support additional planning scenarios, and manual workarounds had become permanent components of the operational process. While each modification solved an immediate business need, the combined result was a system whose complexity far exceeded its original design.

Perhaps the most important observation from the architectural audit was that the workbook itself was never the true system. The spreadsheet merely served as the implementation layer for a much larger operational decision process. Hidden beneath the worksheets was an implicit decision engine responsible for translating inventory data into procurement recommendations through a combination of inventory calculations, safety stock policies, replenishment logic, material classifications, supplier lead times, and buyer-specific business rules.

Recognizing this distinction fundamentally changed the direction of the project. Rather than asking how to automate an Excel workbook, the engineering challenge became identifying, documenting, and reconstructing the operational decision system that the workbook represented. Only after this understanding had been established could the underlying business logic be separated from the spreadsheet and redesigned as a maintainable, centralized software architecture.

PROCUREMENT HEAT MAP

END-TO-END DATA FLOW ARCHITECTURE



Architectural Audit

Motivation

The first engineering decision made during this project was to postpone automation entirely.

Although the existing procurement planning process was widely recognized as labor-intensive, immediately replacing manual steps with code would have risked reproducing years of accumulated complexity. Before designing a new system, it was necessary to understand precisely how the existing system functioned, which components generated value, and which existed only because of historical development.

Rather than viewing the workbook as a spreadsheet, it was treated as a legacy software system requiring architectural analysis.

The objective of the audit was therefore not simply documentation. Its purpose was to identify the functional architecture of the planning system, expose hidden dependencies, quantify technical debt, and establish a foundation upon which a modern decision engine could be designed.

Methodology

Every worksheet within the procurement planning workbook was systematically reviewed and documented.

Each worksheet was evaluated according to a common framework consisting of:

- Primary operational purpose
- Input versus output classification
- Degree of manual intervention
- External and internal data dependencies
- Relationship to other worksheets
- Requirement for future-state architecture
- Opportunities for automation
- Opportunities for elimination or consolidation

Rather than focusing only on calculations, the audit examined how information moved through the system from its initial extraction in SAP to the final purchasing recommendations delivered to buyers.

This approach transformed an undocumented spreadsheet into a traceable information architecture whose individual components could be analyzed independently.

Findings

The architectural audit identified approximately twenty-four worksheets supporting the inventory planning process.

These worksheets performed a diverse range of responsibilities including:

- SAP data ingestion
- Inventory calculations
- Demand aggregation
- Safety stock logic
- Business rule interpretation
- Classification logic
- Output generation
- Dashboard preparation
- Manual validation
- Planner overrides

However, the review also revealed that these responsibilities were not clearly separated.

Many worksheets simultaneously acted as both data repositories and calculation engines. Intermediate calculations were distributed across helper columns, Power Query transformations, nested Excel formulas, lookup tables, and manually maintained configuration sheets. As a result, understanding a single purchasing recommendation often required tracing dependencies across multiple worksheets.

Several worksheets served only as temporary calculation layers or historical artifacts that no longer contributed directly to procurement decisions. Others duplicated information already available elsewhere within the workbook.

The audit therefore demonstrated that the workbook contained significant architectural redundancy despite continuing to produce operationally valuable outputs.

Data Dependency Analysis

One of the most important outcomes of the architectural review was the mapping of data dependencies.

Rather than existing as independent worksheets, the workbook formed a highly interconnected dependency network.

Core inventory information originated from SAP exports, including:

- Purchase Requisitions (ME2N)
- Inventory Balances (MB52)
- Inventory Status Reports (ZMINVENTORY)

These reports were supplemented by manually maintained planning information including:

- Quarterly consumption metrics
- Supplier forecasts
- Business rule tables
- Safety stock parameters
- ABC classification logic
- EOQ logic

Each subsequent worksheet depended upon outputs generated by previous worksheets, creating long calculation chains that ultimately converged within the primary inventory management worksheet responsible for buyer recommendations.

This tightly coupled architecture meant that seemingly minor modifications frequently propagated throughout the workbook, increasing maintenance complexity and reducing confidence in downstream outputs.

Identification of Technical Debt

The architectural audit revealed several forms of accumulated technical debt.

Distributed Business Logic

Critical procurement policies were embedded across hundreds of Excel formulas, Power Query transformations, helper columns, and manually maintained lookup tables rather than being centralized within a single calculation engine.

Redundant Components

Several worksheets contained duplicated calculations, obsolete intermediate outputs, or historical functionality that was no longer required by the planning process.

Manual Data Pipelines

Weekly execution required planners to download SAP reports, copy data into multiple worksheets, refresh calculations in the correct order, validate intermediate outputs, and distribute completed reports manually.

Each step introduced opportunities for human error while increasing total processing time.

Limited Maintainability

Business rule modifications frequently required editing multiple worksheets simultaneously.

Because no centralized calculation layer existed, introducing new planning logic often required updating formulas across numerous tabs, increasing the probability of inconsistencies.

Hidden System Knowledge

Perhaps the greatest source of technical debt was organizational rather than technical.

Much of the system's behavior was understood only through the experience of long-time planners. The workbook itself contained relatively little documentation

explaining why calculations existed or how purchasing recommendations were generated.

As a result, operational knowledge resided primarily within people rather than within the system.

Engineering Implications

The architectural audit fundamentally changed the project's direction.

Initially, the objective had been framed as automating a manual Excel workflow.

However, the audit demonstrated that the spreadsheet was not the problem.

The workbook represented years of accumulated procurement expertise encoded through formulas, planning assumptions, business rules, and operational heuristics. Simply automating existing worksheets would have transferred the same architectural complexity into software without addressing its underlying causes.

Instead, the audit established a new engineering objective:

Design a modular decision system that separates data, business logic, and presentation into independent architectural layers while preserving the procurement knowledge embedded within the legacy system.

This conclusion became the guiding principle for every subsequent design decision.

Rather than recreating spreadsheets in code, the project sought to reconstruct the operational decision system itself.

Design Objectives

The architectural audit demonstrated that the primary limitation of the legacy planning system was not the use of Excel itself, but the manner in which business logic, operational data, and user interaction had become tightly coupled over years of incremental development. Consequently, the objective of this project extended well beyond process automation. The redesign sought to establish a modern decision support architecture capable of supporting procurement operations through greater transparency, maintainability, scalability, and decision consistency.

Rather than replicating the existing workbook in another programming language, the project was guided by a set of engineering principles that prioritized long-term system quality alongside immediate operational improvements.

Separate Business Logic from Presentation

In the legacy system, business rules were embedded throughout Excel formulas, Power Query transformations, lookup tables, and helper worksheets. This architecture made the decision making process difficult to understand, maintain, and modify.

The redesigned system separates these responsibilities into independent layers. Data acquisition, calculation logic, business-rule evaluation, and reporting now operate independently, allowing each component to evolve without requiring changes throughout the remainder of the system.

This separation also improves transparency by making procurement logic explicit rather than hidden within spreadsheet formulas.

Establish a Single Calculation Engine

One of the primary design goals was to eliminate multiple sources of computational logic.

Instead of distributing calculations across dozens of worksheets and formulas, every inventory calculation, demand aggregation, classification, replenishment metric, and purchasing recommendation is executed within a centralized PowerShell calculation engine. This architecture establishes a single source of computational truth while simplifying validation, testing, maintenance, and future enhancements.

Standardize Procurement Decisions

Inventory planners routinely evaluate hundreds or thousands of materials under changing business conditions. While experienced buyers develop strong intuition, relying exclusively on individual interpretation can introduce unnecessary variation into operational decisions.

The redesigned system therefore standardizes procurement recommendations through a centralized business rule engine.

Rather than asking each buyer to independently interpret inventory conditions, the system evaluates reorder point status, inventory position, supplier information, safety stock, lead times, and planning logic to generate consistent recommendations while preserving human oversight for final decision-making.

Reduce Manual Operational Effort

The previous planning process required repeated manual interaction including SAP report extraction, worksheet preparation, Power Query refreshes, validation, formula maintenance, and report distribution.

An important design objective was therefore to reduce repetitive operational work without reducing planner visibility into the underlying calculations.

Automation was applied primarily to deterministic computational tasks, allowing procurement professionals to dedicate more time to supplier management, exception handling, and strategic purchasing decisions rather than spreadsheet maintenance.

Improve Explainability

Operational decision support systems must produce recommendations that users can understand and trust.

For this reason, explainability became a fundamental design requirement.

Each recommendation generated by the Heatmap Engine can be traced through intermediate inventory calculations, safety stock logic, reorder point evaluations, and business rule interpretation. Rather than functioning as a black box optimization model, the system provides transparent reasoning that allows buyers to understand why a particular recommendation was generated.

This transparency simplifies validation, accelerates user adoption, and supports continuous improvement of procurement policies.

Design for Maintainability

Operational planning systems continue evolving as business requirements change.

Accordingly, the architecture was designed so that future enhancements could be implemented without restructuring the underlying calculation engine.

Business-rule tables, planning parameters, classification thresholds, safety stock factors, and configuration logic remain externalized wherever appropriate, allowing planners to update operational policies without modifying application code. Likewise, new calculations, dashboard features, or reporting outputs can be incorporated while preserving the integrity of the core calculation engine.

Build a Foundation for Future Decision Systems

The Procurement Heatmap Engine was intentionally designed as more than an automation project.

By separating operational knowledge from spreadsheet infrastructure, the resulting architecture establishes a reusable framework for future decision-support applications.

Capabilities such as automated purchase order recommendations, supplier prioritization, predictive inventory analytics, demand forecasting, exception reporting, optimization algorithms, and AI assisted procurement planning can now be developed on top of a centralized and well-defined computational foundation rather than being constrained by spreadsheet architecture.

In this sense, the project represents the first stage in the evolution of a broader decision engineering platform.

Design Philosophy

Throughout the redesign process, one principle remained constant:

The objective was not to automate spreadsheets. The objective was to engineer a system that helps people make better procurement decisions.

Every architectural decision, from centralizing calculations to externalizing business rules and integrating interactive dashboards, was evaluated according to a single question:

Does this improve the quality, consistency, transparency, and efficiency of operational decision making?

This principle ultimately shaped every component of the Procurement Heatmap Engine and established the foundation for its continued evolution as an enterprise decision support platform.

System Architecture

The Procurement Heatmap Engine was designed using a layered systems architecture that separates data acquisition, computational logic, business rules, and decision presentation into independent components. This architectural approach reduces system coupling, improves maintainability, and enables future expansion without requiring fundamental redesign.

Rather than embedding calculations throughout spreadsheets, every stage of the planning process performs a single well-defined responsibility. Data flows sequentially through each layer before being transformed into standardized procurement recommendations and operational dashboards.

The resulting architecture establishes a single source of truth for inventory planning while preserving complete traceability from SAP source data to the final buyer recommendation.

Architectural Overview

The system consists of five primary architectural layers:

1. Data Acquisition
2. Data Processing
3. Decision Engine
4. Reporting Layer
5. Operational Distribution

Each layer performs an independent function while exposing standardized outputs to the next stage of the pipeline.

This separation minimizes dependencies between components and allows individual layers to evolve independently without disrupting the overall system.

Layer 1 – Data Acquisition

The first layer is responsible for collecting operational information from enterprise systems.

Weekly SAP reports provide the primary operational inputs required for procurement planning, including inventory balances, purchase orders, future demand, material master information, and standard cost data.

Additional planning information—including safety stock configuration, ABC classification thresholds, EOQ parameters, supplier forecasts, and planning

overrides—is maintained externally through configuration worksheets rather than embedded within application code.

This design ensures that operational policies remain editable by planners while preserving the integrity of the calculation engine.

The complete input dataset includes:

- Inventory Master
- Consumption History
- MRP Demand
- Inventory Position
- Open Orders
- Material Cost
- ABC Logic
- Safety Stock Logic
- EOQ Logic
- Business Rule Configuration

Together, these datasets form the operational foundation upon which all subsequent calculations are performed.

Layer 2 – Data Processing

Once acquired, all operational data enters the centralized PowerShell calculation engine.

This layer performs the complete computational transformation of raw operational information into standardized planning metrics.

Processing responsibilities include:

- Data validation
- Data cleansing
- Aggregation
- Material-level joins
- Demand calculations
- Inventory calculations
- Cost calculations
- Spend calculations
- Lead time calculations
- Material classification
- Safety stock determination
- Reorder point calculations

Unlike the legacy workbook, these calculations are no longer distributed across worksheets or Excel formulas.

Instead, every computational step is executed sequentially within a centralized and version-controlled codebase.

This design dramatically improves maintainability while ensuring that every output can be reproduced consistently from identical input data.

Layer 3 – Decision Engine

After operational metrics have been calculated, the system evaluates procurement business rules.

Rather than presenting buyers with raw inventory data, the Decision Engine interprets inventory conditions using procurement policies developed by the organization.

The engine evaluates variables including:

- Daily demand
- Current inventory
- Inventory under inspection
- Open purchase orders
- Lead times
- Safety stock
- Reorder points
- Material classifications
- Business exceptions

These calculations produce standardized inventory health metrics including:

- ROP Status
- Days to ROP
- Percent Safety Stock Remaining
- Inventory Health
- Required Action

The resulting business-rule engine transforms operational metrics into actionable purchasing recommendations such as:

- Pull In
- Below Reorder Point
- Push Out
- Inventory Healthy
- R&D Material
- Obsolete Material

By centralizing procurement policy within software rather than spreadsheets, the system produces consistent recommendations while preserving complete transparency into the reasoning behind every decision.

Layer 4 – Reporting & Decision Support

Following business-rule evaluation, processed information is delivered to the reporting layer.

Rather than generating static reports, the Procurement Heatmap Engine provides multiple decision-support environments designed for different organizational roles.

These include:

Leadership Dashboard

Provides executive visibility into:

- Inventory health
- Procurement workload
- Inventory spend
- Pull-in opportunities
- Materials below reorder point

Buyer Workbench

Provides buyers with prioritized daily execution lists including:

- Suggested action
- Material
- Supplier
- Inventory position
- Percent safety stock

- Buyer ownership

Inventory Health Analytics

Supports operational investigation through:

- ABC analysis
- Supplier dependency
- Inventory risk
- Action distribution
- Root-cause exploration

Future Inventory Projection

Projects inventory coverage and identifies future shortage risk across multiple planning horizons.

Buyer Workload Dashboard

Measures procurement ownership, workload distribution, supplier concentration, and inventory exposure across purchasing teams.

Layer 5 – Operational Distribution

The final layer delivers information directly to procurement personnel.

Instead of distributing complex Excel workbooks requiring manual interpretation, the Heatmap Engine produces standardized outputs that can be consumed immediately.

Operational outputs include:

- Heatmap Output Workbook

- Power BI Dashboards
- Daily Buyer Reports
- Leadership Inventory Reports
- Standardized Business Rule Recommendations

Because all outputs originate from the same centralized calculation engine, every user observes consistent information regardless of reporting format.

This establishes a true single source of truth for inventory planning.

Architectural Principles

Several engineering principles guided the final architecture.

Modularity

Each architectural layer performs one primary responsibility.

Scalability

New calculations, dashboards, and business rules can be added without redesigning the entire application.

Maintainability

Business logic is centralized, documented, and version controlled.

Transparency

Every recommendation can be traced back through intermediate calculations to its originating SAP data.

Reusability

The architecture provides a reusable foundation for future decision-support systems beyond procurement.

Architectural Summary

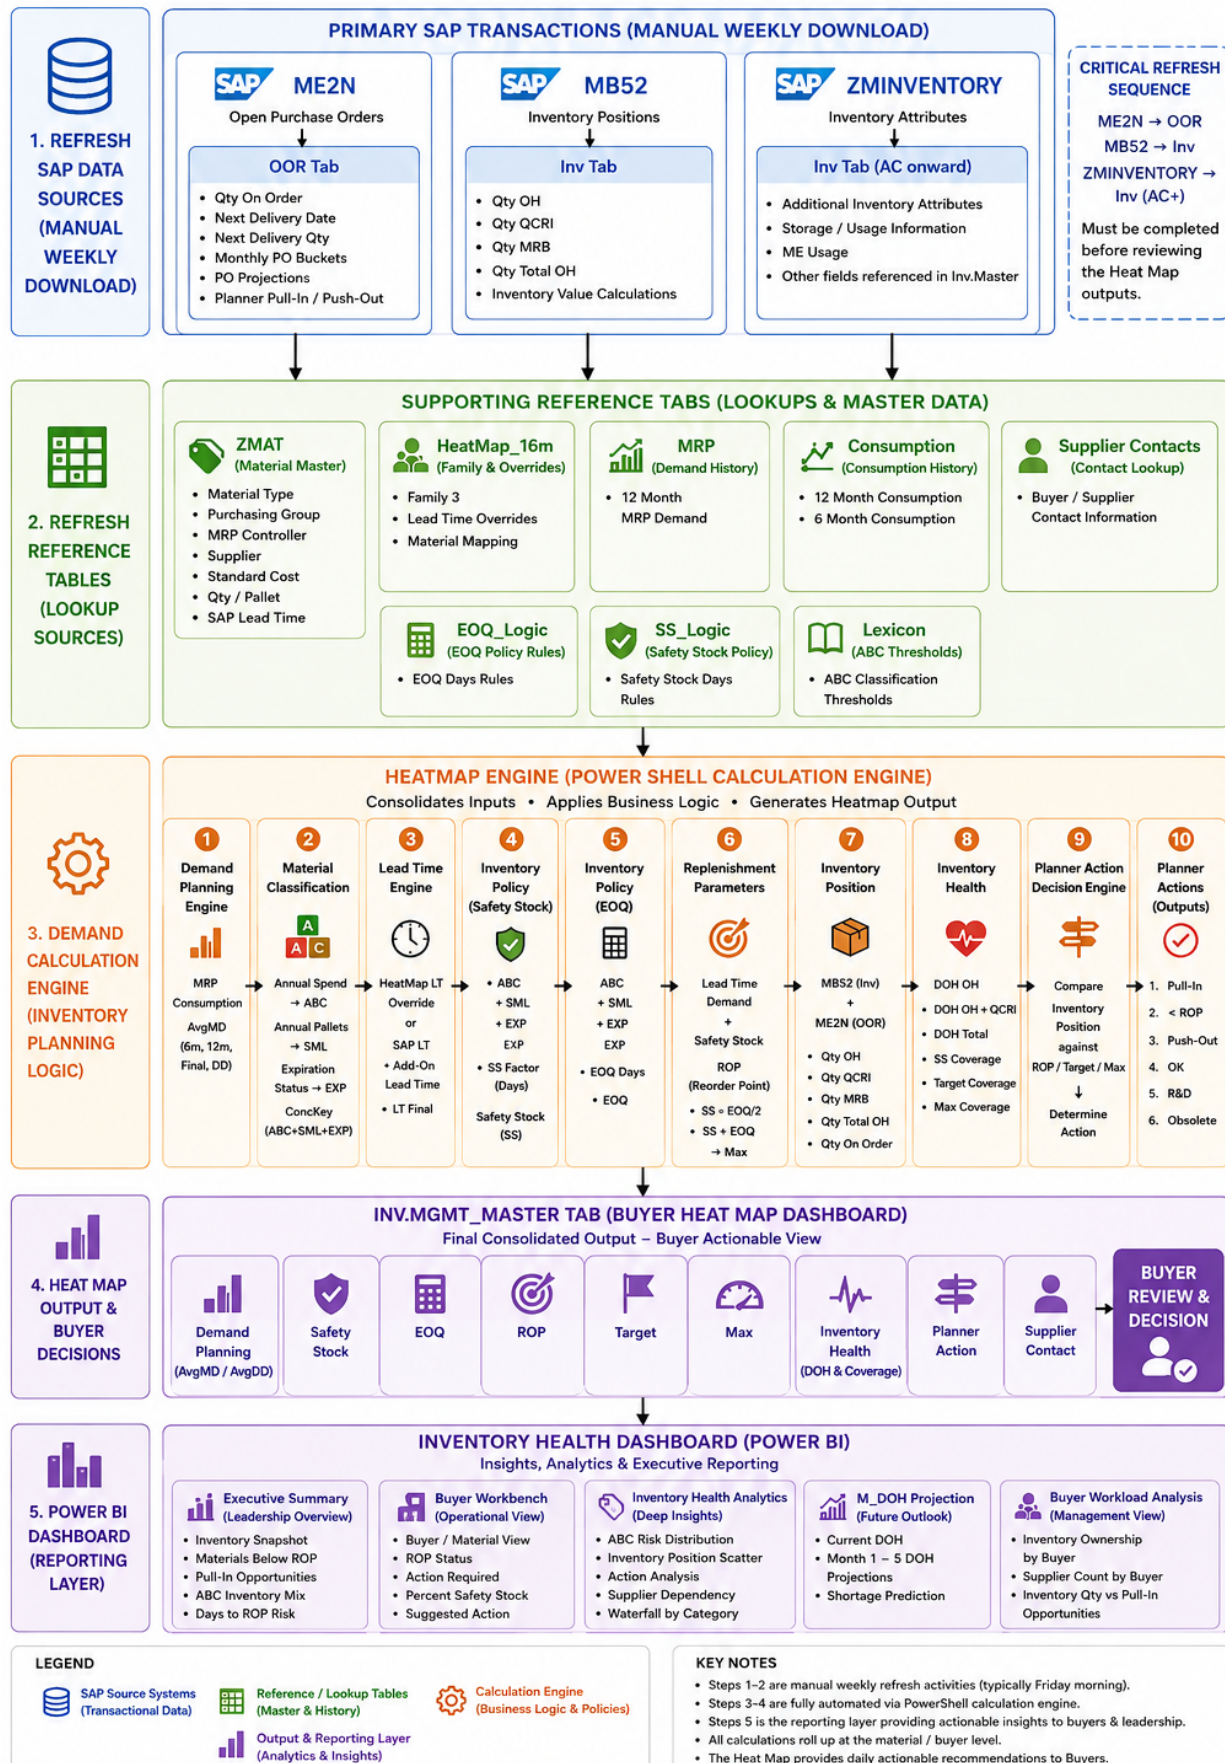
The Procurement Heatmap Engine represents a transition from spreadsheet-based operational planning to software engineered decision support.

Instead of relying on interconnected worksheets, manual calculations, and distributed business logic, the architecture organizes procurement planning into independent computational layers that transform enterprise operational data into transparent, standardized, and explainable purchasing recommendations.

This layered design not only improves the efficiency and reliability of current procurement operations but also establishes the technological foundation for future capabilities including predictive analytics, optimization, intelligent scheduling, automated purchase order recommendations, and AI-assisted operational planning.

HEATMAP ENGINE TO PBI DASHBOARD

END-TO-END DATA FLOW ARCHITECTURE



Engineering the Decision Engine

The Procurement Heatmap Engine was designed to transform raw operational data into standardized procurement decisions through a centralized computational pipeline. Rather than relying on spreadsheet formulas distributed throughout multiple worksheets, every planning calculation is executed sequentially within a single PowerShell-based decision engine.

The primary responsibility of the engine is not merely to calculate inventory metrics, but to convert operational information into explainable purchasing recommendations. Each stage builds upon the outputs of the previous stage, gradually transforming enterprise data into decision-ready information.

Data Integration

The computational process begins by importing operational datasets extracted from SAP together with planning configuration tables maintained by the procurement organization.

These datasets include inventory balances, historical demand, future material requirements, purchase orders, material costs, planning parameters, safety stock configuration, EOQ logic, and business-rule tables.

Because each source contains only one portion of the planning problem, the first responsibility of the engine is to consolidate these independent datasets into a unified material-level view of the supply chain.

Data validation is performed during import to ensure consistency before any calculations begin.

Demand Modeling

Once the datasets have been integrated, the engine estimates expected material demand.

Historical consumption and future MRP demand are aggregated to calculate a consolidated estimate of material usage. Average monthly demand is then converted into average daily demand, creating a common planning metric used throughout the remainder of the engine.

Rather than relying on isolated demand snapshots, this approach combines multiple planning perspectives into a single standardized measure that supports downstream inventory calculations.

Inventory Position

The engine next evaluates the complete inventory position for every material.

Instead of considering only released inventory, the calculation incorporates multiple inventory states, including:

- On-hand inventory
- Inventory under quality inspection
- Material Review Board (MRB) inventory
- Open purchase orders

These values are consolidated into a comprehensive inventory position representing both current and incoming supply.

This integrated inventory view provides a significantly more accurate representation of material availability than evaluating individual inventory categories independently.

Material Classification

Not every material requires the same planning strategy.

Accordingly, each material is classified according to both financial importance and operational characteristics.

Annual material spend is calculated using forecast demand and standard cost before assigning an ABC classification.

Additional classifications—including pallet utilization, expiration characteristics, and planning attributes—are combined into a consolidated classification key used throughout subsequent planning logic.

By incorporating material characteristics into the decision process, the engine adapts planning policies according to business importance rather than applying identical rules across the entire inventory portfolio.

Safety Stock Determination

Safety stock represents one of the most critical planning variables within the procurement process.

The engine determines safety stock using configurable planning factors together with demand history, lead time information, pallet quantities, and organizational planning rules.

Where planner overrides exist, the system preserves those values while continuing to calculate all remaining inventory metrics automatically.

Separating configurable planning policies from computational logic enables procurement teams to modify inventory strategies without requiring changes to application code.

Reorder Point Calculation

Once demand, inventory, lead time, and safety stock have been established, the engine calculates the reorder point for every material.

The reorder point represents the inventory threshold below which replenishment activity should begin.

The calculation considers:

- Average daily demand
- Effective lead time
- Safety stock requirements

From these values the engine derives additional planning metrics including:

- Days to Reorder Point
- Percent Safety Stock Remaining
- Current Inventory Coverage

These metrics provide quantitative measures of inventory health while establishing the foundation for subsequent business-rule evaluation.

Business Rule Evaluation

The computational metrics generated by the engine become meaningful only after they are interpreted through procurement policy.

The Business Rule Engine represents the final reasoning layer within the computational architecture.

Rather than presenting buyers with raw inventory numbers, the engine evaluates inventory conditions using predefined procurement logic.

Each material is assigned a standardized inventory status such as:

- Below Reorder Point
- Healthy Inventory
- Excluded Material

The engine then evaluates the operational action required, including:

- Pull In
- Push Out
- Replenish
- No Action Required
- Research & Development
- Obsolete Material

Finally, these evaluations are translated into human-readable procurement guidance that allows buyers to immediately understand both the inventory condition and the recommended response.

Explainable Decision Making

One of the primary design objectives of the Heatmap Engine was explainability.

Unlike many automated decision-support systems, every recommendation generated by the engine can be traced through each intermediate calculation.

For every material, procurement professionals can understand:

- The underlying demand estimate
- Current inventory position
- Lead time assumptions
- Safety stock calculations
- Reorder point determination
- Business-rule evaluation
- Final purchasing recommendation

This transparency increases user confidence, simplifies validation, and enables continuous refinement of planning policies as operational requirements evolve.

Engineering Summary

The Procurement Heatmap Engine does considerably more than automate spreadsheet calculations.

It transforms disconnected operational datasets into structured procurement knowledge through a deterministic sequence of engineering calculations and business-rule evaluations.

By centralizing every computational step within a single decision engine, the system replaces fragmented spreadsheet logic with an explainable, maintainable, and scalable architecture capable of supporting both current procurement operations and future intelligent decision-support applications.

The result is not simply faster calculations—it is a standardized decision-making framework that enables procurement teams to act consistently, confidently, and proactively across thousands of inventory planning decisions.

Decision Logic: Encoding Procurement Expertise

The most significant engineering challenge addressed by the Procurement Heatmap Engine was not computational complexity—it was knowledge representation.

Procurement planning depends upon a combination of quantitative analysis and operational expertise. Experienced buyers routinely evaluate inventory positions, supplier behavior, lead times, safety stock policies, and business priorities before determining the appropriate purchasing action. Although this reasoning had evolved over many years, much of it remained implicit within spreadsheet formulas or institutional knowledge rather than existing as explicitly documented decision logic.

One of the primary objectives of the Heatmap Engine was therefore to transform procurement expertise into a structured, transparent, and maintainable decision framework.

Rather than replacing human decision making, the engine standardizes the repetitive analytical reasoning that precedes every purchasing decision, allowing buyers to focus their expertise on exceptions, supplier negotiations, and strategic planning.

From Metrics to Decisions

Inventory metrics by themselves rarely produce operational value.

Knowing that inventory equals 4,500 units, safety stock equals 2,000 units, or Days to Reorder Point equals 18 days does not immediately tell a planner what action should be taken.

The Heatmap Engine bridges this gap by interpreting operational metrics through a structured decision framework.

Instead of presenting buyers with isolated numbers, the system converts inventory conditions into standardized procurement recommendations.

This transformation represents the transition from descriptive analytics to operational decision support.

Inventory Health Assessment

Each material is first evaluated according to its overall inventory condition.

The engine determines whether inventory currently exists below the calculated reorder point, remains within acceptable operating levels, or should be excluded from procurement planning due to business-specific conditions such as research materials or obsolete inventory.

This evaluation produces a standardized Reorder Point Status consisting of three possible outcomes:

- 1 – Below Reorder Point
- 2 – Inventory Healthy
- 3 – Excluded

By reducing thousands of inventory records into standardized operational states, the system provides buyers with an immediate understanding of inventory health across the entire material portfolio.

Action Classification

Once inventory health has been established, the engine determines the operational response required.

Unlike traditional inventory reports that require planners to interpret data manually, the Heatmap Engine directly evaluates procurement actions based upon inventory position, future demand, purchase orders, lead times, and planning policies.

Possible actions include:

- **Pull In** — accelerate existing purchase orders to reduce future shortage risk.
- **Below Reorder Point** — initiate replenishment because projected inventory falls below planning thresholds.
- **Push Out** — delay incoming purchase orders to reduce excess inventory.
- **OK** — current inventory conditions require no intervention.
- **Research & Development** — material excluded from normal purchasing activity.
- **Obsolete** — material removed from active planning consideration.

These standardized action categories create a common operational language across procurement teams while ensuring that identical inventory conditions generate identical recommendations.

Business Rule Interpretation

Although inventory status and recommended action provide significant operational guidance, procurement decisions frequently require additional context.

To address this need, the engine combines Reorder Point Status and Action into a unified Business Rule Interpretation.

This layer translates computational outputs into human-readable recommendations that buyers can immediately understand and act upon.

Rather than requiring users to mentally combine multiple planning variables, the engine performs this reasoning automatically before presenting the recommendation.

The resulting output functions less as a report and more as an operational decision assistant.

Explainable Procurement Decisions

One of the defining characteristics of the Decision Engine is explainability.

Every recommendation generated by the system can be traced through a deterministic sequence of intermediate calculations.

For every material, buyers can understand:

- The calculated demand.
- The current inventory position.
- Safety stock requirements.
- Lead time assumptions.
- Reorder point calculations.
- Inventory health assessment.
- Business rule evaluation.
- Final procurement recommendation.

This transparency distinguishes the Heatmap Engine from black-box decision-support systems.

Recommendations are not merely generated—they are fully explainable.

This characteristic increases user confidence, simplifies validation, and supports continuous refinement of procurement policies as organizational requirements evolve.

Standardizing Operational Knowledge

Perhaps the most important contribution of the Business Rule Engine is organizational rather than computational.

Before the redesign, much of the procurement process depended upon individual buyer experience.

Two planners reviewing identical inventory conditions could reasonably arrive at different conclusions based upon personal interpretation.

By explicitly encoding procurement policies into software, the Heatmap Engine creates a consistent decision framework that reduces unnecessary variation while preserving human judgment for complex or exceptional situations.

Rather than replacing procurement professionals, the system amplifies their expertise by ensuring that routine analytical reasoning is performed consistently across every planning cycle.

Engineering Summary

The Decision Engine represents the point at which operational data becomes organizational knowledge.

Demand, inventory, supplier information, lead times, and planning policies are no longer viewed as isolated metrics. Instead, they are synthesized into standardized recommendations that support faster, more transparent, and more consistent procurement decisions.

In doing so, the Procurement Heatmap Engine shifts the role of technology from reporting historical information to actively supporting operational decision making.

This transition, from calculating inventory metrics to engineering explainable decisions, represents the central contribution of the project and establishes a reusable framework for future intelligent decision-support systems.

Decision Support & Operational Intelligence

The computational outputs produced by the Procurement Heatmap Engine achieve operational value only when they support effective human decision making. While the Decision Engine standardizes inventory calculations and

procurement recommendations, the reporting layer transforms these outputs into actionable operational intelligence tailored to different levels of the procurement organization.

Rather than producing static reports, the Heatmap Engine delivers information through a series of interactive Power BI dashboards designed around specific decision-making responsibilities. Each dashboard serves a distinct user group while maintaining a common computational foundation. This architecture ensures that executives, procurement managers, and buyers all operate from the same underlying data while receiving information appropriate to their operational needs.

The reporting layer therefore functions not as a visualization platform, but as the primary interface between the computational decision engine and the procurement organization.

Design Philosophy

Traditional reporting systems emphasize historical visibility.

The objective of the Procurement Heatmap Engine was fundamentally different.

Every dashboard was designed to answer one operational question:

What decision should I make next?

Rather than maximizing the number of charts or metrics displayed, each visualization was selected according to its ability to support a specific procurement decision.

Throughout the reporting layer, emphasis was placed on:

- Decision prioritization
- Operational transparency
- Risk identification
- Executive visibility
- Buyer accountability

- Explainable recommendations

This philosophy transformed the dashboard from a passive reporting tool into an active operational workspace.

Leadership Inventory Health Overview

The Leadership Inventory Health Overview provides executives and procurement leadership with a real-time assessment of overall inventory performance.

Rather than reviewing multiple operational reports independently, leadership receives a consolidated view of procurement health through key performance indicators including:

- Total Inventory Quantity
- Total Inventory Spend
- Materials Below Reorder Point
- Pull-In Opportunities
- Inventory at Risk
- Buyer Ownership
- Material Count

Visualizations such as ABC inventory distribution, Days to Reorder Point analysis, and inventory health summaries enable leadership to identify systemic risks before they develop into operational disruptions.

Instead of monitoring isolated transactions, management gains visibility into the overall health of the procurement system.

Buyer Workbench

While leadership focuses on organizational performance, buyers require operational execution tools.

The Buyer Workbench functions as the daily operational interface for procurement planners.

Each buyer receives a prioritized list of materials requiring attention together with the information necessary to act immediately.

For every material, the dashboard presents:

- Material Number
- Buyer
- Supplier
- Current Inventory
- Inventory Health Status
- Recommended Action
- Percent Safety Stock Remaining
- Business Rule Interpretation

Instead of manually reviewing hundreds of materials, buyers immediately understand which items require intervention and why.

This significantly reduces manual analysis while improving consistency across procurement activities.

Inventory Health Analytics

Procurement decisions frequently require investigation beyond individual materials.

The Inventory Health Analytics dashboard supports root-cause analysis by allowing users to explore inventory performance from multiple operational perspectives.

Interactive visualizations include:

- ABC Risk Distribution
- Supplier Dependency Analysis
- Inventory Position Scatterplots
- Action Distribution
- Inventory Health Waterfall

These visualizations enable procurement teams to identify patterns that would be difficult to detect through spreadsheet analysis alone.

Examples include:

- Supplier concentration risk
- Inventory imbalance
- Financial exposure
- High-risk material categories
- Procurement workload trends

By supporting exploratory analysis, the dashboard enables procurement professionals to move beyond reactive decision making toward continuous operational improvement.

Inventory Projection

Inventory planning requires anticipating future conditions rather than responding only to current inventory levels.

The Inventory Projection dashboard estimates future inventory coverage across multiple planning horizons using current inventory positions, demand forecasts, and replenishment assumptions.

Projected Days on Hand enables buyers to identify shortages before they occur, providing additional time for supplier communication and purchasing adjustments.

This forward-looking capability supports:

- Proactive procurement planning
- Reduced expediting costs
- Improved supplier collaboration
- Increased supply continuity

Rather than reacting to shortages after they occur, procurement teams can intervene while operational flexibility still exists.

Buyer Workload Analysis

Effective procurement depends not only on inventory health but also on organizational capacity.

The Buyer Workload dashboard provides management with visibility into buyer ownership, supplier assignments, inventory responsibility, and workload distribution.

Key visualizations include:

- Inventory Ownership by Buyer
- Supplier Count by Buyer
- Pull-In Opportunities by Buyer
- Inventory Exposure by Buyer

This information enables procurement leadership to balance workloads, identify resource constraints, and support capacity planning across the purchasing organization.

Operational decisions therefore extend beyond inventory management to include organizational effectiveness.

From Reporting to Decision Support

One of the primary design goals of the reporting layer was to move beyond traditional business intelligence.

Conventional dashboards frequently answer descriptive questions such as:

- What happened?
- How much inventory do we have?
- Which supplier delivered late?

The Procurement Heatmap Engine instead emphasizes prescriptive decision support by answering questions including:

- Which materials require immediate attention?
- Which purchase orders should be accelerated?

- Where is inventory at greatest risk?
- Which buyers require additional support?
- Which suppliers present the greatest operational exposure?

This transition represents a shift from reporting operational performance toward actively improving operational decision making.

Operational Impact

The reporting layer completes the transformation initiated by the Decision Engine.

Raw SAP data is progressively refined into operational metrics, inventory health assessments, procurement recommendations, and finally role-specific decision-support environments.

Because every dashboard originates from the same centralized calculation engine, executives, managers, and buyers operate from a single source of truth while receiving information tailored to their specific responsibilities.

The result is a procurement organization capable of making faster, more consistent, and more transparent decisions while reducing manual analysis, improving collaboration, and increasing organizational visibility across the entire inventory planning process.

Rather than serving as the end of the system, the reporting layer represents the point at which engineered decision intelligence becomes actionable operational behavior.

Results & Operational Impact

The Procurement Heatmap Engine transformed the procurement planning process from a workbook-centered workflow into a centralized decision-support system. While the underlying inventory policies remained consistent with existing procurement practices, the manner in which those policies were executed changed substantially.

The redesign produced improvements across four dimensions: system architecture, operational execution, decision quality, and organizational visibility.

Architectural Transformation

The most significant outcome of the project was architectural rather than computational.

The legacy planning system relied on a large Excel workbook that had evolved over several years into an interconnected collection of worksheets, Power Query transformations, formulas, lookup tables, helper columns, and manually maintained logic.

The redesigned architecture centralized these responsibilities within a single PowerShell-based decision engine.

Instead of distributing calculations throughout the workbook, inventory logic now exists within a version-controlled software application capable of producing deterministic and repeatable outputs.

This transformation established several long-term engineering benefits:

- A single calculation engine.
- A single source of truth for inventory planning.
- Centralized business-rule management.
- Modular software architecture.
- Improved maintainability.
- Simplified future enhancement.

Perhaps most importantly, the redesign separated operational knowledge from spreadsheet implementation, allowing procurement policy to evolve independently of the computational platform.

Simplification of the Legacy Planning System

The architectural audit revealed that the legacy workbook contained approximately twenty-four worksheets supporting various aspects of inventory planning.

A detailed review demonstrated that several worksheets no longer contributed meaningful operational value, while others performed intermediary calculations or duplicated functionality available elsewhere within the workbook.

Rather than recreating every worksheet in software, the redesign focused on preserving business functionality while eliminating unnecessary architectural complexity.

This simplification reduced technical debt and produced a significantly cleaner computational architecture without sacrificing planning capability.

Instead of viewing every worksheet as a required system component, the project distinguished between operational requirements and historical implementation choices.

That distinction became one of the most important engineering outcomes of the redesign.

Operational Efficiency

The redesigned workflow substantially reduced manual interaction throughout the weekly procurement planning cycle.

Previously, buyers were required to:

- Export multiple SAP reports.
- Populate numerous workbook tabs.
- Execute Power Query refreshes.
- Validate intermediate calculations.
- Navigate multiple worksheets.

- Interpret inventory conditions manually.
- Prioritize purchasing actions individually.

Following implementation of the Heatmap Engine, these computational activities became centralized within an automated processing pipeline.

The procurement team's effort shifted away from preparing data and toward evaluating recommendations, collaborating with suppliers, and resolving operational exceptions.

Rather than replacing procurement expertise, the system redirected buyer attention toward higher-value decision making.

Standardization of Procurement Decisions

Prior to the redesign, purchasing recommendations depended heavily upon manual interpretation of inventory metrics.

Although experienced buyers generally reached appropriate conclusions, identical inventory conditions could occasionally be interpreted differently depending upon planner experience or analytical approach.

The Business Rule Engine introduced a standardized decision framework that evaluates inventory health using consistent computational logic.

Every material is now assessed according to the same planning methodology before generating a procurement recommendation.

This standardization improves consistency across buyers while preserving human judgment for situations requiring contextual knowledge beyond the available operational data.

Organizational Visibility

One of the most valuable outcomes of the project was the creation of organization-wide inventory visibility.

Previously, information existed primarily within operational workbooks designed for individual buyers.

The introduction of interactive Power BI dashboards extended this visibility throughout the procurement organization.

Leadership now has immediate access to:

- Overall inventory health.
- Procurement workload.
- Inventory spend.
- Supplier exposure.
- Inventory risk.
- Future shortage projections.

At the operational level, buyers receive prioritized worklists based upon standardized business rules rather than manually reviewing every material independently.

This alignment enables executives and procurement planners to make decisions using the same underlying operational information while viewing it through interfaces tailored to their respective responsibilities.

Foundation for Future Decision Systems

Although the Procurement Heatmap Engine was developed to support inventory planning, the architecture extends well beyond a single procurement application.

By separating data ingestion, computational logic, business rules, and visualization into independent layers, the project established a reusable framework for future decision-support systems.

Potential extensions include:

- Automated purchase order recommendations.
- Predictive inventory forecasting.

- Supplier prioritization.
- Intelligent scheduling.
- Optimization algorithms.
- AI-assisted procurement planning.
- Exception management.
- Digital supply chain analytics.

Rather than representing the end of a modernization effort, the Heatmap Engine provides the technological foundation upon which future operational intelligence capabilities can be constructed.

Engineering Reflection

Perhaps the most important outcome of this project was a shift in perspective.

The original objective appeared to be automating a complex spreadsheet.

However, the engineering process revealed that the spreadsheet itself was never the true system.

The real system consisted of the procurement knowledge embedded within years of operational decision making.

By documenting that knowledge, separating it from spreadsheet infrastructure, and reconstructing it as a modular software architecture, the project transformed a legacy planning workbook into an explainable operational decision engine.

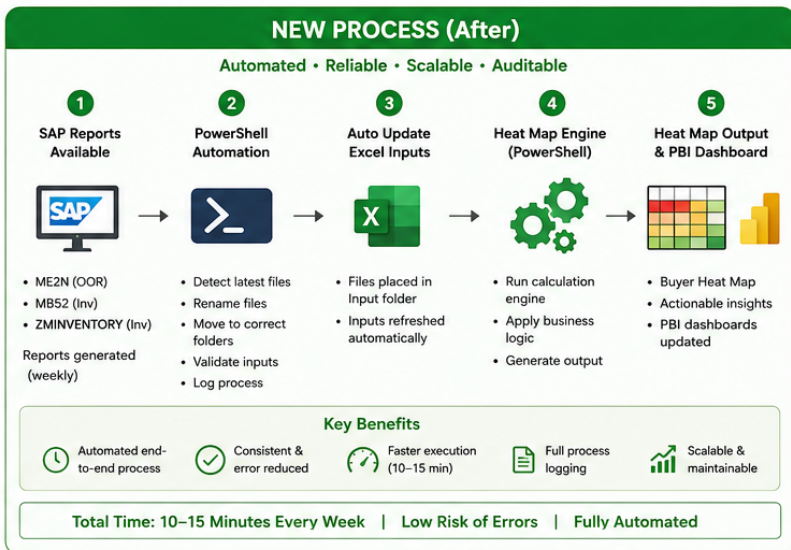
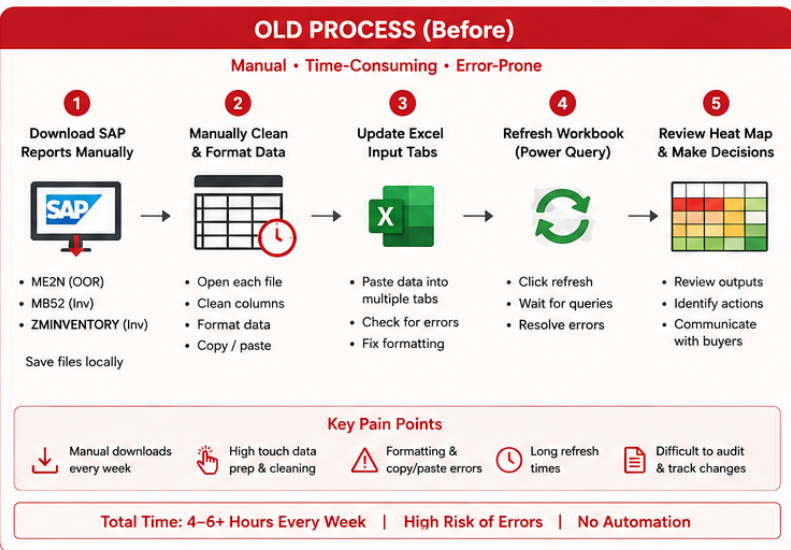
The result is more than an automation solution.

It is an example of how systems engineering principles can modernize operational processes by making organizational knowledge explicit, maintainable, scalable, and reusable.

In doing so, the Procurement Heatmap Engine demonstrates that the greatest value of digital transformation lies not merely in replacing manual work, but in engineering better decision systems.

HEAT MAP PROCESS: OLD vs NEW

From Manual, Excel-Driven Workflow to Automated, End-to-End Solution



PROCESS COMPARISON SUMMARY

Aspect	OLD PROCESS	NEW PROCESS	IMPACT
Data Acquisition	Manual download of 3 SAP reports	Automated file detection & management via PowerShell	Eliminates manual weekly work
Data Preparation	Manual cleaning, formatting, copy/paste	Automated file placement & input refresh	Improves data quality & consistency
Processing / Calculations	Power Query refresh (multiple steps)	PowerShell calculation engine (business logic applied)	Faster processing & more reliable
Output Generation	Excel heat map only	Excel output + Power BI dashboards	Richer insights & visualization
Time Required	4-6+ hours per week	10-15 minutes per week	~90% time savings
Error Risk	High (manual handling & copy/paste)	Low (automated & validated)	More reliable decisions
Audit & Traceability	Difficult to track changes	Process logging & file tracking	Full audit trail & transparency
Scalability	Hard to scale / add new data	Easy to scale & maintain	Supports growth & future needs

BOTTOM LINE

The new process transforms the Heat Map from a manual, error-prone Excel workbook into an automated, reliable, and scalable inventory planning solution.

RESULT:
Better Data. Faster Insights. Stronger Decisions.



Engineering Reflections

Although the Procurement Heatmap Engine ultimately produced a centralized decision-support platform, the most valuable outcomes of the project were not technical. Throughout the redesign process, a series of broader engineering principles emerged that fundamentally changed how operational systems should be understood, analyzed, and redesigned.

Rather than viewing this project simply as an automation effort, it became an exercise in systems engineering, organizational learning, and decision architecture.

The following reflections summarize the most significant lessons that emerged during the development process.

The Spreadsheet Was Never the System

At the beginning of the project, the workbook appeared to be the problem.

Its numerous worksheets, distributed formulas, Power Query transformations, and manual refresh steps created the impression that replacing Excel would solve the underlying challenges.

However, the architectural audit demonstrated that the spreadsheet itself was only the visible implementation of a much larger operational system.

The true system consisted of procurement policies, planning assumptions, inventory strategies, supplier knowledge, and organizational decision-making processes that had gradually accumulated over many years.

The spreadsheet merely served as the medium through which those decisions were expressed.

This realization fundamentally changed the engineering objective.

The project was no longer about replacing Excel.

It became an effort to understand and reconstruct the operational decision system hidden beneath it.

Automation Without Understanding Creates Better Complexity

One of the strongest temptations in automation projects is to reproduce existing workflows in software.

Doing so often accelerates execution while preserving the same architectural limitations that existed previously.

This project demonstrated that automation should never precede understanding.

Only after every worksheet, dependency, calculation, and business rule had been documented was it possible to distinguish operational requirements from historical implementation choices.

The resulting software therefore automates procurement logic rather than automating spreadsheets.

This distinction proved essential to producing a maintainable architecture.

Business Rules Are Engineering Assets

Throughout the legacy planning process, procurement expertise existed in many forms.

Some knowledge resided in Excel formulas.

Some appeared within lookup tables.

Some existed only through the experience of individual planners.

Initially these elements appeared to be isolated operational details.

Over time it became clear that they collectively represented the organization's decision-making framework.

Once these business rules were documented explicitly, they became engineering assets.

Rather than remaining embedded within spreadsheets, organizational knowledge became software that could be validated, maintained, improved, and reused across future decision-support applications.

Explainability Builds Trust

Many modern decision-support systems emphasize increasingly sophisticated predictive models.

While predictive accuracy remains important, operational adoption depends equally upon transparency.

Procurement professionals must understand why recommendations are generated before they are willing to rely upon them.

For this reason, explainability became a deliberate design objective rather than a secondary feature.

Every recommendation produced by the Heatmap Engine can be traced through demand calculations, inventory metrics, safety stock logic, reorder point evaluation, and business-rule interpretation.

This transparency transformed the system from an automated calculator into a trusted operational partner.

The Objective Was Better Decisions, Not Better Dashboards

Perhaps the most important lesson emerged late in the project.

Initially, success appeared to depend upon producing better reports and more attractive dashboards.

However, the dashboard itself was never the product.

The product was improved operational decision making.

Every architectural decision—from centralized calculations to standardized business rules and role-specific dashboards—was ultimately evaluated according to one question:

Does this help procurement professionals make better decisions?

This shift in perspective transformed the project from a reporting initiative into a decision engineering initiative.

Engineering Organizational Knowledge

One unexpected realization was that organizations possess enormous quantities of operational knowledge that often remain undocumented.

Experienced employees routinely make excellent decisions because they understand relationships that have never been formally represented within software.

By documenting these relationships and translating them into transparent computational logic, engineering becomes a mechanism for preserving institutional knowledge rather than merely automating repetitive work.

This principle extends far beyond procurement.

Many organizational processes contain valuable decision frameworks that remain hidden within spreadsheets, emails, procedures, and individual experience.

Engineering those systems begins not with programming, but with understanding.

Reflection

Looking back, the Procurement Heatmap Engine was never simply a software project.

It became an exercise in discovering how operational knowledge flows through an organization, how that knowledge becomes encoded within technology, and how engineering can make those decision processes more transparent, maintainable, and scalable.

The most valuable outcome was not replacing formulas with code.

It was learning that every operational system tells a story about how an organization makes decisions.

Understanding that story is the first responsibility of the engineer.

Only then can better systems be designed.

In many ways, this project reshaped my own understanding of engineering.

I no longer see engineering primarily as the creation of software or automation.

I see it as the discipline of designing systems that enable people to make better decisions.

The Procurement Heatmap Engine represents the first realization of that philosophy, but it also established the foundation for a broader vision: applying systems engineering, artificial intelligence, and operations research to improve human decision making across increasingly complex organizational environments.

Toward Intelligent Decision Systems

The Procurement Heatmap Engine demonstrates that operational knowledge can be transformed into a centralized, explainable, and maintainable decision-support system. Although the project focused specifically on procurement planning, the underlying engineering principles extend far beyond inventory management.

At its core, the system illustrates a broader idea: organizations generate enormous amounts of operational data, but their greatest challenge is rarely the availability of information. Instead, the challenge lies in transforming that information into decisions that are timely, consistent, explainable, and actionable.

The Procurement Heatmap Engine represents one implementation of that philosophy.

Rather than viewing this project as a finished product, it should be understood as the first step toward a broader family of intelligent decision systems capable of supporting increasingly complex organizational environments.

Predictive Decision Support

The current system evaluates inventory conditions using deterministic business rules derived from procurement policies.

A natural next step is to extend the platform with predictive capabilities.

Demand forecasting models could anticipate future inventory shortages before reorder points are reached, while predictive supplier analytics could estimate delivery risk using historical performance data.

These capabilities would enable procurement teams to move beyond responding to current inventory conditions toward proactively managing future operational uncertainty.

The Decision Engine would therefore evolve from answering:

"What should we do today?"

to

"What should we prepare for tomorrow?"

Optimization

Many procurement decisions involve balancing competing objectives.

Reducing inventory may increase stockout risk.

Increasing safety stock improves resilience but ties up working capital.

Accelerating purchase orders improves material availability while increasing inventory carrying costs.

These tradeoffs naturally lend themselves to optimization techniques drawn from operations research.

Future versions of the Heatmap Engine could incorporate mathematical optimization models that recommend purchasing strategies based on organizational objectives such as:

- Minimizing inventory cost.
- Maximizing service level.
- Reducing supplier risk.
- Improving inventory turnover.
- Balancing buyer workload.

Rather than replacing procurement expertise, optimization would provide buyers with quantitatively supported alternatives for evaluating complex planning decisions.

Artificial Intelligence

The modular architecture developed in this project also establishes a foundation for integrating artificial intelligence into procurement planning.

Large language models and machine learning systems should not replace established operational policies.

Instead, they can augment human decision making by providing additional context, identifying unusual patterns, summarizing operational conditions, and supporting investigation of complex supply chain scenarios.

Potential applications include:

- Intelligent exception summaries.
- Automated supplier risk assessments.
- Natural-language explanations of procurement recommendations.
- AI-assisted inventory investigations.
- Conversational interfaces for operational analytics.

Within this framework, artificial intelligence functions as a collaborative decision-support technology rather than an autonomous decision maker.

Decision Engineering Beyond Procurement

Although developed for procurement planning, the underlying architecture is intentionally domain independent.

The same engineering approach can be applied wherever organizations must transform complex operational data into structured decisions.

Potential applications include:

- Financial performance management.
- Budget variance analysis.
- Manufacturing operations.
- Capacity planning.

- Workforce scheduling.
- Quality management.
- Healthcare operations.
- Project portfolio management.

Across each of these domains, the engineering challenge remains fundamentally the same:

Transform organizational knowledge into explainable decision-support systems.

The Procurement Heatmap Engine therefore serves not only as a procurement solution, but also as a reusable architectural pattern for operational decision engineering.

A Personal Engineering Philosophy

Throughout the development of this project, one realization became increasingly clear.

Engineering is often associated with building software, writing algorithms, or designing technical infrastructure.

While those activities remain important, this project suggests a broader perspective.

Engineering can also be understood as the discipline of designing systems that help people make better decisions.

The Procurement Heatmap Engine did not replace procurement professionals.

It reduced repetitive analytical work, standardized routine reasoning, and created greater transparency into the decision-making process.

The expertise, judgment, and experience of buyers remain essential.

Technology simply enables those capabilities to be applied more effectively.

This perspective has shaped my own approach to engineering.

Rather than viewing automation as the ultimate objective, I increasingly view technology as a mechanism for augmenting human judgment through systems that are transparent, explainable, and designed around the needs of decision makers.

Conclusion

The Procurement Heatmap Engine began as an effort to modernize a complex spreadsheet-based planning process.

Through architectural analysis, systems engineering, and software redesign, the project ultimately became something much larger.

It demonstrated that operational decision systems can be reconstructed from legacy infrastructure without sacrificing the organizational knowledge embedded within existing processes.

By separating data, computation, business rules, and decision support into independent architectural layers, the project established a scalable platform capable of supporting future analytics, optimization, and artificial intelligence applications.

More importantly, it demonstrated that the greatest opportunity for engineering does not lie solely in processing more data.

It lies in helping people make better decisions with the information they already possess.

As organizations continue generating increasingly complex operational data, the role of engineers will extend beyond building software toward designing intelligent systems that transform information into understanding and understanding into action.

The Procurement Heatmap Engine represents one step toward that future.