Agentic AI Assistant for Supplier Management

by

Jing Liu

and

John Zhang

Submitted on May, 2026 in Partial Fulfillment of the
Requirements for the Degree of Master of Applied Science in Supply Chain Management
at the Massachusetts Institute of Technology

ABSTRACT

This capstone project explores the application of agentic AI to improve supplier segmentation at a global healthcare company managing a complex global supplier network. The current segmentation process relies on manual, binary yes/no assessments across 13 questions, producing outcomes that are often subjective, inconsistent across category managers, and difficult to trace back to supporting evidence. To address these challenges, we designed and developed a proof-of-concept Supplier Management Assistant powered by a multi-agent AI system. The solution integrates retrieval-augmented generation (RAG) with a two-tier data strategy, combining question-mapped primary sources with broader contextual knowledge, to automatically retrieve, synthesize, and present structured evidence for each segmentation decision. A human-in-the-loop design ensures that category managers retain final decision authority while benefiting from AI-generated recommendations grounded in existing enterprise data platforms. Expected outcomes include a reduction in segmentation time, improved decision consistency and traceability, and scalable coverage across the full supplier base. Working with the stakeholders, the project also delivers a suggested phased roadmap for scaling the solution from proof-of-concept through pilot to enterprise-wide adoption, positioning the Sponsor company to transition from manual, resource-constrained supplier governance toward an intelligence-driven, AI-augmented supplier management capability.

Capstone Advisor: Dr. Thomas Koch

Title: Postdoctoral Associate

TABLE OF CONTENTS

1. INTRODUCTION

This project was conducted in collaboration with a global healthcare organization (the "Sponsor"). The Sponsor operates across multiple healthcare-related business areas and maintains a broad international presence across numerous countries and operational sites worldwide. As a large multinational organization with a complex global supplier network, the Sponsor faces growing needs for scalable, consistent, and data-driven supplier management processes.

The Sponsor's business is divided into two main segments, pharmaceuticals and medical technology, covering a full spectrum of healthcare solutions. The pharmaceutical branch of the company, focused on the development of novel drugs to help cure and treat patients. The medical technology business unit encompasses a plethora of medical technologies and treatments used during surgical procedures.

As a global leader with complex procurement, supplier network, and rigorous supplier management ecosystem, the Sponsor provides an ideal context for exploring leveraging agentic AI system in their supplier management processes.

1.1. Motivation

The Sponsor has a supplier base that corresponds to its size and scale. Despite a world-class Supplier Management (SM) framework laid out in their SM playbook, the Sponsor faces challenges in terms of effectively executing, scaling and sustaining the desired best practices from the framework.

According to the playbook, the Sponsor's SM process relies on segmentation as the critical "step zero" that informs all downstream governance decisions, including performance reviews, risk management, and supplier engagement strategies. However, today this segmentation process is highly manual, inconsistent across teams, and dependent on fragmented evidence scattered across multiple systems. As a result, valuable data is underused, decision quality varies by region and individual, and meaningful supplier insights are often delayed or lost.

At the same time, the Sponsor's broader vision is to evolve toward an integrated, intelligence-driven Supplier Management ecosystem. Achieving this requires a clear path forward: not only solving today's segmentation pain points but also defining how Agentic AI could orchestrate future SM workflows across personas, data streams, and functional handoffs.

Therefore, the Sponsor is motivated to explore how GenAI and agent-team architectures can (1) immediately strengthen segmentation quality and consistency through an experimental Supplier Management (segment & engage) Assistant, which serves as an experimental Proof-of-Concept (POC) to surface tangible value today and (2) outline a scalable multi-year blueprint for AI-enabled supplier governance. By doing so, the Sponsor aims to strengthen objective-data-backed decision making, improve traceability, increase operational efficiency, and unlock the strategic value hidden in its supplier data, ultimately building a more reliable and future-ready SM capability.

1.2. Problem Statement

Supplier segmentation at the Sponsor plays a critical role in guiding supplier management strategy, including how suppliers are prioritized, monitored, engaged, and governed across the organization's global operations. Segmentation outcomes influence downstream activities such as relationship management, risk oversight, innovation collaboration, and allocation of internal resources. However, despite its strategic importance, the current segmentation process remains highly subjective, manual, and dependent on individual interpretation rather than objective evidence. Category managers are required to classify suppliers based on 13 segmentation questions, covering innovation, performance, and risk aspects of suppliers, but each question is answered only with a binary yes/no. There is no centralized system capturing the underlying drivers behind the yes/no answer, such as KPIs, performance metrics, financial indicators, or risk data.

Instead, category managers must piece together insights from fragmented, inconsistent, and occasionally unverified sources, e.g. finance data, spend reports, and ad-hoc conversations with relevant stakeholders, which may also vary widely by category. As a result, different managers may answer the same question differently for the same supplier, creating inconsistency, lack of traceability, and limited transparency in segmentation outcomes. Category managers repeatedly expressed that the process feels “not always objective” and “sometimes difficult to justify with concrete evidence.”

This issue is amplified by the dynamic nature of supply chains, where a supplier’s financial health, innovation capability, or risk exposure may shift rapidly. Today, identifying such changes relies almost entirely on human effort, for example, subscribing to newsletters, manually reading industry reports, checking financial documents, or monitoring risk events. In addition, because of limited time capacity, category managers cannot realistically monitor all relevant information. These manual checks are incomplete, inconsistent, and often delayed, meaning segmentation decisions frequently do not reflect up-to-date supplier realities.

At the same time, the Sponsor already owns multiple data platforms - a financial data source (filings and financial data, news, competitive intelligence, patents reports on the supplier level), market intelligence tool (industry, category and market intelligence), and risk tracker (risk event data on the supplier level) - that contain exactly the type of evidence needed to support segmentation. However, these sources remain largely under-utilized because the data is difficult to navigate, scattered across systems, or presented in long, unstructured documents that managers do not have time to review.

The combination of subjective binary decisions, lack of evidence capture, dynamic supplier conditions, and under-utilized data sources leads to inconsistent segmentation, inefficient processes, and missed opportunities to leverage supplier insights for strategic decision-making.

Therefore, the Sponsor needs a solution that (1) brings structure, evidence, and objectivity to today’s segmentation workflow – which is what the project team is designing as a working POC / prototype named as Supplier Management Assistant (segment and engage) backed by a multi-agent AI system, and (2) defines a scalable blueprint for how Agentic AI could orchestrate supplier segmentation and overall governance end-to-end in the future.

1.3. Key Questions

This project seeks to address the following questions around how agentic AI can improve supplier segmentation at the Sponsor:

1. How can agentic AI transform the current subjective, binary yes/no segmentation into a more structured, evidence-driven process that captures the reasoning behind each answer?
2. How can a multi-agent system process and summarize evidence from the Sponsor's existing data sources to support more consistent and defensible segmentation decisions?
3. Given that relevant data sources vary widely by category, how can the system accommodate different evidence types while maintaining a consistent decision framework?
4. How should the system be designed to keep human judgment central to the final segmentation decision, while still reducing manual effort?
5. What would a realistic future-state roadmap look like for scaling this approach beyond the current proof-of-concept?

1.4. Project Goals

The project aims to deliver a working proof-of-concept and a forward-looking roadmap. Specifically, the goals are to:

- Design and develop an experimental Supplier Management Assistant that guides category managers through a more structured, evidence-driven segmentation process, reducing subjectivity in the current binary approach.
- Demonstrate that a multi-agent AI system can process pre-loaded supplier data from the Sponsor's existing platforms and generate evidence-backed recommendations for each segmentation question.
- Establish a standardized decision-logic framework that defines how evidence is interpreted, validated, and presented to category managers to ensure transparency and reproducibility.
- Outline a phased roadmap for how this approach could evolve toward automated data integration, broader supplier coverage, and extended governance workflows in future phases.

2. STATE OF THE PRACTICE

The main objective of this capstone project is to improve the quality, consistency, and objectivity of supplier segmentation at the Sponsor through the application of AI technologies. The current segmentation process is still largely manual, with numerous touchpoints required from the category managers. Our solution focuses on utilizing agentic AI and retrieval-augmented generation (RAG) to automate the segmentation process, and improve accuracy, useability, and scalability to the business. The research below dives into specific methods that were applied to the final solution. The first being rule-based segmentation, which is the foundational logic which the AI layer basis decision-making off. Secondly, retrieval-augmented generation (RAG) is a specific process which allows for data processing. Finally, multi-agentic AI systems is the specific framework used to link together the various functionalities needed to process the wide range of data sources. The end product will be an autonomous system comprised of multiple AI agents, able to assist category managers throughout the segmentation process.

2.1. Rule-Based Segmentation

The current segmentation process at the Sponsor is structured as a rule-based, binary scoring system. A rule-based approach is a logic-driven method that utilizes predefined “if-then” statements to produce consistent outputs. In this specific context, category managers may provide binary (Yes/No) responses to qualitative and quantitative criteria. The system will then aggregate totals per category and quantify the resulting supplier segmentation.

While this approach may appear simplistic at first, utilizing a rule-based approach can offer distinct advantages. The main benefit is a clear criterion, which allows a previously variable process to have a defined answer. The Sponsor has over 35,000 suppliers, therefore having a process that is consistent and repeatable is vital to a robust supplier network. A rule-based approach allows for the optimization of the segmentation process, often 40-60% faster than more complex, weighted score-based approaches (Samandari et al. 2013). The rule-based logic reduces the risk of human-related errors and ensures the consistent classification of suppliers.

However, rule-based segmentation poses the risk of oversimplification of complex business environments. Once parameters are set, especially on a binary scale, it does not allow a system to be adaptable to any potential changes (Tursunalieva et al. 2024). The inflexibility posed by this simplification can result in inaccurate segmentations based on broad targets. For example, a supplier that is slightly below a pre-set benchmark may be grouped with a supplier that is severely below the benchmark. Thus, creating areas of ambiguity for category managers, resulting in missed opportunities and misaligned planning for suppliers that are incorrectly segmented. To mitigate these risks, organizations can continuously update parameters through periodic reviews.

The Sponsor already has an existing schedule which it uses to evaluate supplier performance throughout the year. However, introducing an agentic system allows this process to be iterated more frequently within the same time period. This allows for segmentation results to be more accurate as inputs utilize more recent data. Thus, by keeping the current segmentation logic, we can ensure scalability and reliability without requiring a revamp of the existing process or retraining of existing stakeholders.

2.2. Retrieval-Augmented Generation (RAG)

When scoping for this project, it quickly became apparent that the Sponsor has robust existing sources of data. It is critical that the solution that we implemented can leverage this data source. A common critique of current AI models is the risk of "hallucinations"- instances where the system produces incorrect information. Errors are caused by mismatches between the model's statistical predictions and the correct information. Due to large language model's (LLM) learning via existing data patterns, it may be unable to deliver an accurate recommendation with limited parametric memory consisting of the original data training set. A retrieval-augmented generation (RAG) system allows for the agentic system to access and ingest the Sponsor's proprietary data that it would not normally have access to, allowing recommendations to be accurate and relevant to the company.

Implementing a RAG system allows for an improvement in data integrity and the accuracy of the agent's responses. When compared to a traditional LLM, the RAG framework offers some distinct advantages, including deep contextual awareness, source citation, and fewer hallucinations (Kimothi, 2024). These factors allow the agentic system to utilize recent, internal data to make recommendations, all without requiring the retraining of the existing LLM. The Sponsor's current segmentation process leans heavily on the in-house expertise of category managers, and decision traceability and transparency are difficult to manage. Using the RAG framework allows the agentic system to replicate this industry-specific knowledge by having access to proprietary data such as contracts, spend data, and company dashboard. The introduction of the RAG system allows for a more structured process as well as opportunities to assist human users with ambiguous segmentation decisions.

However, certain drawbacks are associated with RAG systems. A major risk is the LLM utilizing misleading or incorrect data to make inaccurate decisions. Many of the presented advantages of RAG were based on the assumption of a "clean retrieval" environment, essentially that the LLM has access to correct and clean data (Zeng et al., 2025). This is not necessarily the case with the current company data sources, which are fragmented across various sources- each with different formats and various update frequencies. Additionally, the size of the dataset can have negative effects on both accuracy and response time. Given the scale of the Sponsor's operations, this is an important variable to consider, as each category manager can be responsible for over 100 suppliers in each review cycle. Thus, as we narrow down the data sources utilized, it is important for these drawbacks to be fully factored into the decision process.

2.3. Multi-Agent Systems

Given the scope and complexity of the current segmentation process, we began to explore solutions that could sufficiently cover the needed data processing capabilities while maintaining efficiency and useability. Multi-agent systems allow users to leverage the efficiencies of individual AI agents, while also utilizing the synergies between agents to fully cover the complex task that is presented. Within a multi-agent system, individual AI agents act as autonomous entities, each tasked with unique roles within the system. Each output is then orchestrated and synthesized to produce an output that matches the requested criteria from the user.

A key functionality of the segmentation solution requires the AI agents to be able to extract, evaluate, and synthesize data from a variety of sources. This is also a major challenge within the project, because as even though we plan to leverage an RAG system to allow agents to make segmentation decisions, it requires an accurate and sufficiently sized dataset. Coordinating this data stream is an entire process in itself, as scoping conversations with current category managers revealed that existing data sources that are available for analysis are scattered and heterogeneous.

There are multi-agentic frameworks which demonstrate the distribution of responsibilities within a system. Figure 1 details an example of multi-agentic flow that details a knowledge discovery process, starting from data selection and ending with interpretation/evaluation. Data needs to be collected from a variety of sources from internal company repositories to external data sources. Additionally, it come in a variety of file formats and need to be processed for the segmentation agent to properly review the data. Leveraging a similar flow to summarize data sources is a critical responsibility for the segmentation agentic system.

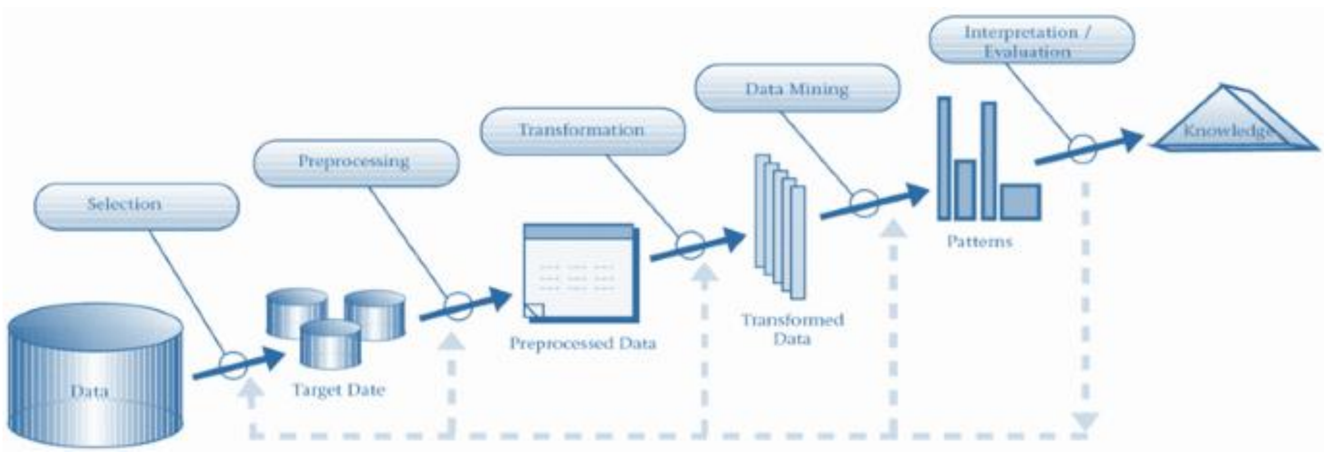

Figure 1. The Knowledge Discovery Process (Ait-Mlouk et al. 2016)

AI agents can play a plethora of roles within a system. Users can customize and make additions based on desired functions. Some examples of agentic roles include (Jaleel et al. 2020):

- Collaborative Agents- agents that are autonomous and feature social behaviors with other agents
- Interface Agents- agents that operate within the autonomous software space and can perform end-user tasks
- Information Agents- agents that can manipulate data from various sources
- Hybrid Agents- agents that incorporates multiple agent roles/functionalities
- Reactive Agents- agents which operates within a constrained environment, its acts are limited to agent environment
- Cognitive Agents- agents which aims to find solutions and execute, its able to process learning and act to make decisions.

Identifying the tasks that are necessary for the segmentation process is critical when designing a multi-agentic system. For this project, information agents are needed to compile and process various data sources. Collaborative agents are then needed to summarize and combine findings from the information agents. Finally, cognitive agents must be able to understand and propose outputs based on data inputs. This hierarchy of agents must work together to collect, process, and provide feedback from a variety of data sources. By breaking down the segmentation process into a series of steps, each handled by an agent with a specific role, we can leverage the tactical expertise of the AI agent while also mitigating the limitations of an agent operating in isolation within a limited scope. Coordinating multiple agents is vital in ensuring the most comprehensive outcomes, allowing this solution to be scalable to new supplier categories, as the multi-agent approach allows for continuous improvements and learning with each iteration.

3. METHODOLOGY

This methodology section outlines the structured approach used to design, develop, and implement the multi-agentic supplier segmentation tool. It explains how the project is organized into distinct phases: scoping, construction, and review. The technical architecture is detailed, displaying the process of the multi-agent framework, data integration strategy, and human-in-the-loop validation process. Additionally, the end-to-end user experience and workflow are described, illustrating how category managers interact with the system from initial login until final segmentation output. Finally, a dedicated section outlines the data collection and preparation approach, including how primary and contextual data are incorporated.

3.1. Technical Approach and System Architecture

The proof-of-concept is built using LangFlow (orchestration layer) and Streamlit (user interface), and structured around a multi-agent architecture that enables dynamic retrieval, synthesis, and evaluation of supplier-related data to support segmentation decisions.

At a high level, the system follows a modular, orchestrated workflow. The user initiates the process through the Streamlit interface by selecting supplier identifiers (e.g., name and category) and optionally uploading contextual documents. These inputs are processed and passed to the orchestration layer, where an Orchestrator Agent coordinates the end-to-end execution.

The architecture integrates heterogeneous data sources, including structured internal datasets (e.g., spend data, contract data, and dashboards queried via SQL) as well as unstructured external or user-provided documents (e.g., PDFs and PowerPoint files), which are processed through an Azure Document Intelligence pipeline to extract usable text. This enables the system to combine primary data (question-specific) with contextual complementary data for richer reasoning. The core processing logic is distributed across several specialized agents:

- The Data Retrieval Agents (e.g., spend analysis and SQL research) independently query relevant sources in parallel.
- A Compiler Agent aggregates and structures the retrieved evidence into a unified intermediate representation.
- Domain-Specific Agents (e.g., Innovation, Performance, and Risk Agents) interpret the compiled evidence against segmentation criteria.
- A Recommendation Agent applies predefined segmentation rules and synthesizes a final recommendation.

The Orchestrator Agent manages the sequencing and interaction between these components, ensuring a coherent flow from data retrieval to final output. Importantly, the system incorporates a human-in-the-loop mechanism, where users can review, validate, and override intermediate outputs, improving both transparency and decision reliability.

This architecture (Figure 2) enables parallel data processing, traceable evidence generation, and scalable segmentation across a broad supplier base, addressing key limitations of manual and single-agent approaches.

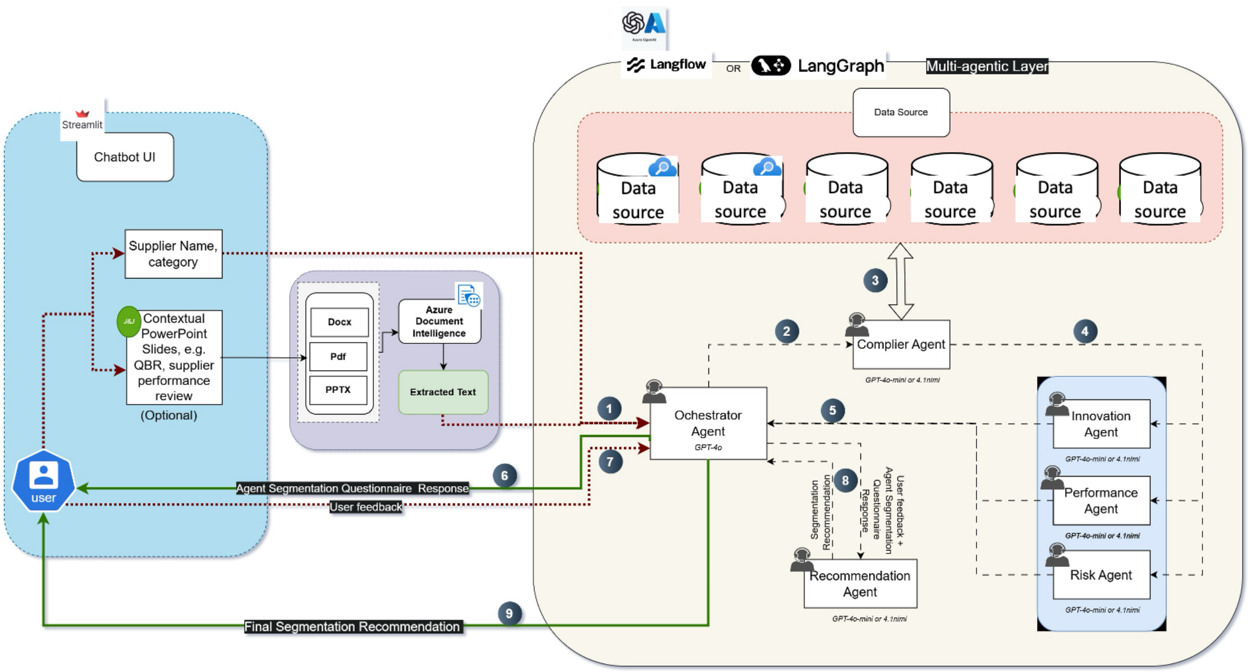

Figure 2. Architecture Diagram for Agentic AI Solution

3.2. Agentic-AI-Enabled Segmentation Workflow

The workflow structures supplier segmentation into a streamlined, end-to-end process where AI aggregates and translates multi-source data into decision-ready insights, while keeping humans in the loop for validation. The process flow is shown in Figure 3.

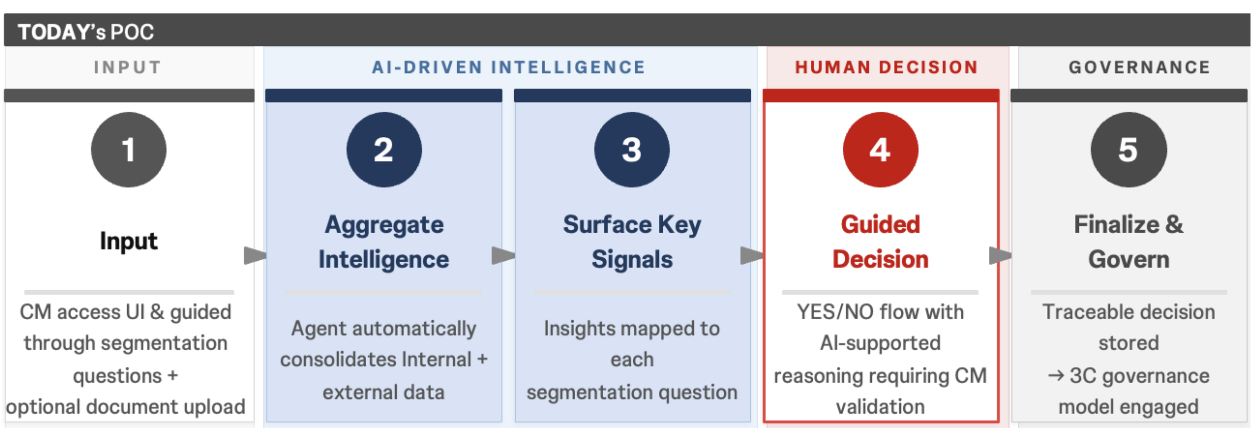

Figure 3. New Process Enabled by Agentic AI Solution

To begin the segmentation process, the user first enters their name and credentials. Upon successful login, the user is prompted to select the correct region, category, and specific supplier. This allows the segmentation tool to correctly identify and pull in the supporting evidence needed.

The specialists agents with access to relevant data sources are then prompted to begin research. First, the spend analysis agent will retrieve the corresponding spend data which is pre-loaded into the AI system. Parameters of reference are also available; additionally, the user can customize ranges as desired. The Structured Query Language(SQL) research agent will also run in tandem with the first agent, reviewing existing company dashboards for data sources that apply to the segmentation question. Once this initial research phase is complete, the compiler AI agent will be prompted to consolidate all data inputs into a final answer. The orchestrator agent will then synthesize inputs from both the specialist agents and the compiler agent to deliver a single flow of information to the recommendation agent.

When all data has either been extracted or uploaded, the recommendation agent will give a summary answer for the final recommendation along with a recommended “Yes” or “No” for segmentation questions. However, the category manager will be able to review and change the agent’s original answer if needed. In this step, the user can upload any relevant supporting documents that the agentic system may not have had access to in the first iteration. The compiler agent will then extract relevant data points and deliver the results to the user. The agent will repeat this process for all 13 questions. Once the survey is complete, the agent will compile the results and deliver the final recommendation. An output dataset will be generated to allow the user to conduct further analysis if desired. Furthermore, the segmentation process will be logged, allowing for reviews in future audits and traceability.

The team is in the review and refinement phase for this model, and we anticipate that there will be discrepancies with the agentic AI’s accuracy and usability depending on the category of supplier that is supporting. During our interviews with current category managers, we observed various dynamics and challenges unique to their respective categories. Challenges include limited public news data, lack of clear criteria definitions, and scattered internal data sources. As we complete the construction phase, knowledge of the nuances of each category will be vital in ensuring that the agent AI is finetuned to meet the needs of category managers and truly be an asset within the supplier segmentation process.

3.3. Data Collection and Preparation

The effectiveness of the agentic supplier segmentation system depends critically on the quality, coverage, and traceability of the underlying data. During the scoping phase, we worked closely with the Sponsor company, process owners, and key stakeholders to identify the most authoritative (“best-practice”) data sources available for each of the 13 segmentation questions. However, a consistent finding across interviews was the absence of a universal “golden source” for every question. As a result, we intentionally designed the data strategy to accommodate uneven data availability while maintaining decision traceability.

To address this reality, the data architecture was structured around two complementary data pathways: (1) question or segment (innovation, performance and risk) level mapped sources and (2) general contextual knowledge sources processed via LLM reasoning.

3.3.1. Two-Tier Data Strategy

Tier 1: Question or Segment Level Mapped Sources

To maximize precision and auditability for high-confidence signals, all of the reliable, high-signal data sources that we identified were explicitly mapped to individual segmentation questions or segments in collaboration with the Sponsor company’s stakeholders. This mapping improves retrieval precision and reduces the risk of irrelevant evidence being surfaced by the model. Table 1 summarizes the primary segmentation dimensions used by the Sponsor and provides representative examples of the data sources typically referenced to support supplier evaluation across innovation, performance, and risk categories.

Segmentation Dimension	Purpose in Segmentation Process	Representative Data Sources
Innovation	Assess supplier innovation capability, collaboration potential, and strategic value creation opportunities	Contracts and agreements, supplier presentations, QBRs, partnership reviews, market intelligence reports
Performance	Evaluate supplier operational performance, continuous improvement capability, sustainability practices, and supplier management metrics	Supplier performance reviews, scorecards, sustainability assessments, supplier diversity dashboards, operational KPIs
Risk	Assess supplier criticality, supply continuity exposure, market concentration, regulatory exposure, and reputational risk	Financial and market intelligence reports, supplier risk monitoring tools, contract data, spend analytics, regulatory and industry reports

Table 1. Supplier Segmentation Dimensions and Representative Data Sources

Tier 2: General Contextual Knowledge Sources

In the meantime, to maximize data coverage and surface weak or emerging signals, broader enterprise and external documents are ingested into the RAG knowledge base. The documents include SEC filings, quarterly business review (QBR) slide decks, supplier performance review slide decks, and other narrative supplier materials that could be uploaded by category managers through the user interface. These sources are not pre-mapped to specific questions. Instead, the LLM performs semantic evidence discovery, identifying relevant signals and mapping them probabilistically to the segmentation question and criteria.

It is worth noting that the data strategy employs a two-tier approach in which Tier 1 leverages stakeholder-validated, question-mapped primary sources to generate high-confidence Yes/No assessments, while Tier 2 uses broader contextual corpora to supplement evidence and surface signals where direct mappings are limited. The fused question-level outputs are then processed through the rule-based segmentation logic from the Sponsor company to produce the final supplier segmentation recommendation. Importantly, the agentic system does not replace the rule logic. Rather, it automates the evidence gathering and binary assessment that feeds the existing governance model. This design choice preserves organizational trust while improving speed and consistency. (See Figure 4.)

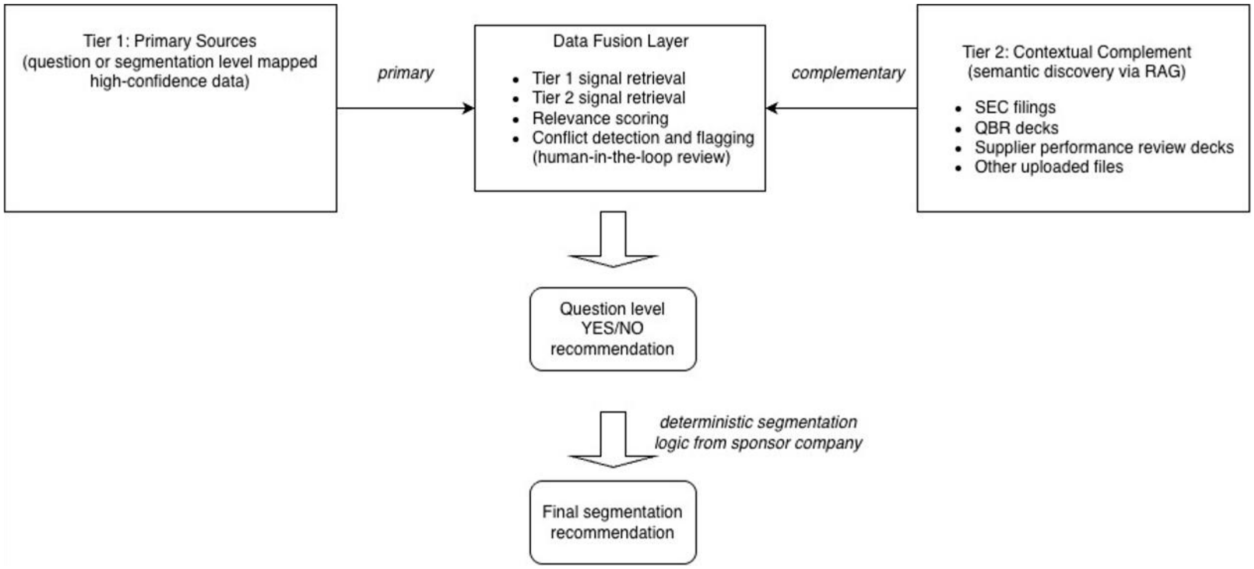

Figure 4. Two-Tier Data Fusion Architecture

3.3.2. Multi-Lens Data Classification

To ensure consistent handling of heterogeneous inputs, all data sources are classified using a unified multi-lens framework that captures origin, structure, and analytical role in the segmentation process. Table 2 summarizes the classification schema applied across the data pipeline.

Data Source	**Origin**	**Structure Type**	**Signal Role in Segmentation**	**Typical Usage**
Sustainability and compliance assessments	Internal	Semi-structured	Direct binary evidence	Primary (Tier 1)
Supplier diversity dashboard	Internal	Structured	Direct binary evidence	Primary (Tier 1)
Spend and contract data	Internal	Structured numerical	Direct binary evidence	Primary (Tier 1)
SmartCube reports	Internal	Structured / semi-structured	Direct binary evidence	Primary (Tier 1)
SEC filings	External	Unstructured and tables	Contextual signal	Secondary (Tier 2)
QBR slide decks	Internal	Unstructured	Contextual signal	Secondary (Tier 2)
Supplier performance review slide decks	Internal	Unstructured	Contextual signal	Secondary (Tier 2)
Other uploaded materials	Internal	Unstructured	Contextual signal	Secondary (Tier 2)

Table 2. Multi-Lens Data Classification

3.3.3. Expected Impact on Segmentation Quality

By combining targeted high-precision data mapping, broad contextual retrieval, multi-modal evidence processing, and rule-based governance, the proposed architecture aims to reduce manual data-gathering burden; improve evidence traceability; surface weak signals earlier; and increase segmentation consistency across category managers. However, overall accuracy will remain partially dependent on the quality and completeness of upstream data, which continues to vary by supplier category.

3.4. Methodology Application

The agentic system follows a specific data pathway to ensure that generated recommendations are accurate and scalable. This section details the technical features that allow the agentic system to utilize AI to carry out the intended outcome. It specifically describes how data is ingested, transformed, retrieved, and evaluated through a combination of RAG, multi-agent orchestration, and rule-based logic. By integrating all these components, the agentic system can extract relevant evidence for final segmentation decisions.

The first substantive step is data ingestion and representation. After the user has entered the correct credentials and uploaded any additional documents to supplement the existing database, the agentic system utilizes a specific technique known as chunking (breaking down large documents into ingestible segments that an LLM can process). The chunking logic used in this system is primarily structural; for example one chunk per spreadsheet row, PDF page, or slide. This helps to ensure that the agent is capturing enough meaningful evidence that is congruent with what a category manager would interpret.

The second step is to retrieve relevant data based on supplier and category. Using the user's inputs for both features, the system will use a query to identify relevant chunks that pertain to the selected supplier. Chunks are ranked by embedding similarly when the vector search is enabled. The number of retrieved results is governed by the top-k parameter, which limits the ranked text segments returned by the query. This ensures that the selected data is relevant and filters the outside noise from low-quality or unrelated content. The retrieved data set forms the qualitative evidence layer; from which the segmentation process will be ultimately based.

The third step involves per-question suggestions which can be changed through human revision. There are three specific specialist agents: innovation, performance, and risk, which are each responsible for analyzing the data sources to identify evidence per each specific category. Upon selecting and analyzing the correct data source, the system then converts retrieved evidence into an initial binary recommendation using heuristic rules. (Let $Yi \in \{0,1\}$ represent the final answer to question i, where 1 indicates the criterion is met and 0 indicates it is not. Let Ai denote the model's suggested answer and Hi the human override. The final outcome is defined as $Yi = Hi$ if a reviewer provides an override, and $Yi = Ai$ otherwise.) This step is critical in ensuring that human judgment remains the final gatekeeper for segmentation decisions.

Finally, the resulting data is processed through the established segmentation rules that the Sponsor's category managers currently utilize. First, a supervisor agent consolidates the inputs from the three specialist agents. Then the compiler agent takes this input and compiles it with the data source processing results from the RAG agent. The final Boolean outcomes are evaluated over the ordered, discrete decision process to determine the final segmentation outcome (Growth, Emerging, Essential, or Transact). The recommendation agent will take all inputs and create a final answer to be provided to the user. This framework ensures reproducibility and alignment with the established governance framework. Upon successful segmentation, the agentic system will create a final summary view, where key metrics such as number of human overrides, audit timestamps, and exportable results are available for the end user. Ultimately, this stage prioritizes interpretability and traceability, allowing the user to have access to the final segmentation decision and the option to explore the tactical decisions taken to achieve this recommendation.

In summary, the proposed methodology integrates a multi-agent, RAG-enabled architecture with an existing rule-based governance framework to support supplier segmentation in a transparent and scalable manner. This approach combines targeted data mapping with broader contextual retrieval to ensure that both precision and coverage are available in evidence collection. Overall, the objective of these methodological steps is to enhance decision quality and velocity while preserving human managerial oversight and established governance.

4. RESULTS AND DISCUSSION

Given the time constraints of the capstone, the project team prioritized the testing and refinement of the solution to ensure a robust and scalable foundation. As a result, comprehensive empirical validation with category managers could not be conducted at scale within the project timeline. However, this approach enables the delivery of a more mature and adaptable solution, positioning the Sponsor to carry out systematic user testing and validation in subsequent phases. Accordingly, based on prototype performance, stakeholder interviews, and alignment with the Sponsor's strategic priorities, this section outlines the expected results and anticipated benefits of the proposed solution.

First, the proposed agentic AI solution is expected to significantly improve process efficiency. By automating data retrieval, aggregation, and initial assessment, the system reduces the manual effort required by category managers. Figure 5 depicts a preview of the user interface that a category manager views. The category manager is able to view existing data sources and upload additional files. All sources are automatically categorized into specific segments that allow the multi-agent system to review and generate an accurate recommendation. Early directional estimates suggest the potential to reduce segmentation time by more than 50%, enabling a substantial increase in the number of suppliers that can be evaluated within each review cycle.

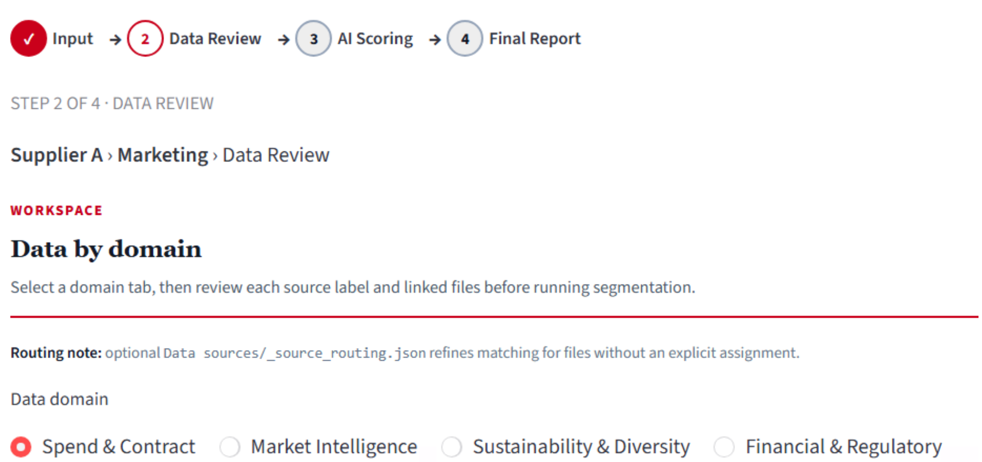

Figure 5. Data Review From Proof-of-Concept Model

Second, the solution is expected to enhance decision quality and consistency. By grounding each segmentation decision in structured evidence retrieved through the RAG framework, the system reduces variability caused by individual interpretation. This is particularly important in the current state, where execution of the same rule-based framework differs across category managers. The introduction of AI-supported reasoning provides a more standardized decision context, leading to more consistent segmentation outcomes across categories.

Third, the system strengthens governance and traceability. Each segmentation decision is accompanied by documented evidence and rationale, creating a fully auditable decision trail. Figure 6 details a sample output from the POC model. The multi-agentic system highlights key findings from relevant documents and converts the RAG vectors into readable findings for the category manager. The synthesis completed by the system is then validated by the category manager and creates a record of segmentation logic that can be accessed for future iterations. This directly addresses one of the key limitations identified in the current process, namely, the lack of transparency and reusability of decision logic. Over time, this capability is expected to support stronger alignment between supplier segmentation and enterprise-level governance frameworks.

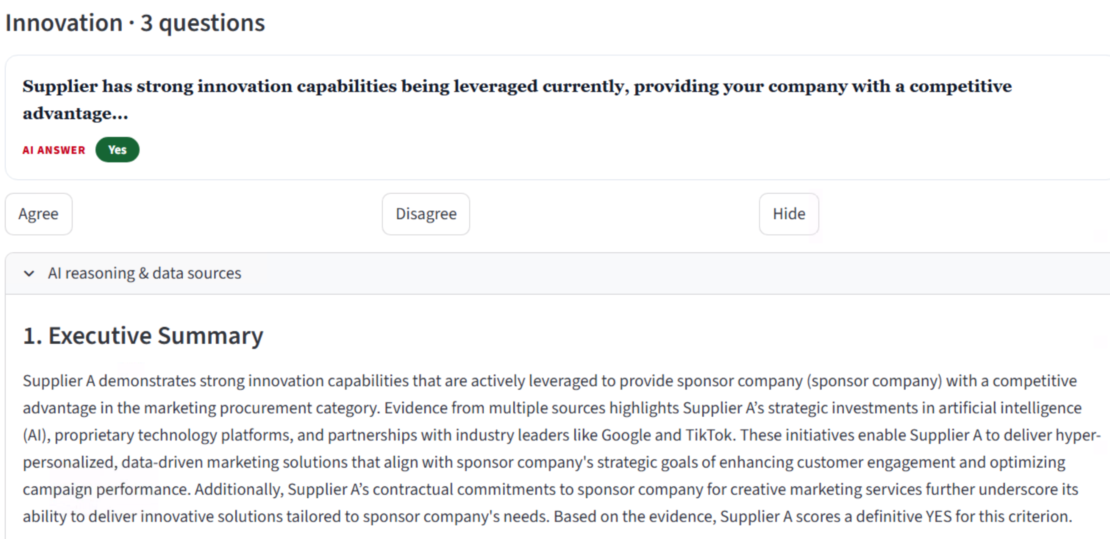

Figure 6. Sample AI Recommendation Output From Proof-of-Concept Model

Finally, the solution establishes a scalable foundation for supplier management transformation. Figure 7 displays the final confirmed recommendation view that the category manager will have the final sign-off on. This approval process will be repeated for each segmentation that the multi-agentic system runs. By enabling segmentation across a significantly larger portion of the supplier base, the system allows organizations to move from selective, resource-constrained segmentation toward a more comprehensive and continuously updated view of supplier criticality. In the current state, category managers appropriately prioritize Growth and Emerging suppliers (those that are strategic or have near-term potential to drive competitive advantage for the business). However, this leaves a large share of transactional, or tail-spend suppliers insufficiently assessed, creating the risk that potentially critical suppliers remain unidentified. By automating evidence of aggregation and initial assessment, the multi-agentic AI solution enables scalable coverage across the full supplier base, ensuring no supplier is overlooked while keeping segmentation continuously up to date.

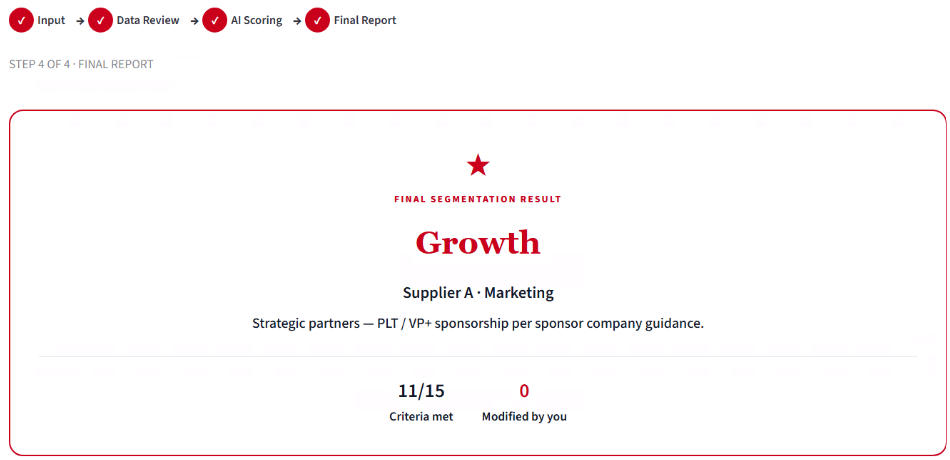

Figure 7. Final Recommendation View From Proof-of-Concept Model

It is important to note that these expected benefits remain partially dependent on upstream data quality and availability, which vary across supplier categories. As such, realized performance may differ across use cases and will require iterative refinement during deployment.

5. RECOMMENDATIONS

Based on the findings from the POC development and stakeholder interviews conducted throughout the project, this section outlines the recommendations for how the Sponsor can build on the current prototype to scale agentic AI within supplier management. The first priority is establishing a phased roadmap that transitions the solution from its current POC state toward broader organizational adoption.

5.1. Future Roadmap and Scaling Strategy

While the current project focuses on developing a POC, the long-term vision extends toward a fully integrated, enterprise-scale supplier management system. The roadmap follows a phased approach: POC → Pilot → Scale-up.

In the POC phase (2026)**,** the primary objective is to validate the feasibility of the agentic AI approach within a controlled environment. This includes deploying the prototype to a limited number of categories and suppliers, with partial data integration and standalone system functionality. At this stage, the focus is on demonstrating directional value, validating the multi-agent architecture, and identifying key limitations related to data availability, model accuracy, and user interaction.

The subsequent pilot phase (2027) aims to expand both functionality and adoption. This phase will involve deeper integration with enterprise data systems, including spend, contract, and supplier intelligence platforms. The system is expected to achieve measurable improvements in segmentation relevance (e.g., increased business acceptance of AI recommendations) and further reductions in manual effort. Additionally, outputs from the segmentation process will begin to feed downstream processes such as category strategy development and supplier governance workflows.

In the scale-up phase (2028 and beyond)**,** the vision is to establish a fully operational agentic AI assistant embedded within the broader procurement ecosystem. At this stage, the system will support end-to-end supplier segmentation across the entire supplier base, with high levels of accuracy and minimal manual intervention. Advanced capabilities may include prescriptive recommendations, automated execution of low-risk operational tasks, and simulation-based “what-if” analysis to support strategic decision-making.

Across all phases, several critical enablers will determine successful scaling. These include improving data quality and standardization, establishing clear performance metrics and “ground truth” definitions, and implementing continuous model monitoring to mitigate drift. Equally important is the adoption of a human-in-the-loop governance model to ensure that AI recommendations remain aligned with business context and stakeholder expectations.

Overall, the roadmap reflects a gradual transition from augmentation to partial automation, where AI first enhances human decision-making and eventually enables a more autonomous, scalable supplier management capability.

REFERENCES

Ait-Mlouk, A., Gharnati, F., & Agouti, T. (2016). Multi-agent-based modeling for extracting relevant association rules using a multi-criteria analysis approach. *Vietnam Journal of Computer Science*, *3*(4), 235–245

Jaleel, H. Q., Stephan, J. J., & Naji, S. A. (2020). Multi-agent systems: A review study. *Ibn Al-Haitham Journal for Pure and Applied Science*, 33(3), 2483.

Kimothi, A. (2024). *A simple guide to retrieval-augmented generation (video edition)*. Manning Publications.

Samandari, H., Walsh, J., & Yueh, E. (2013, July 1). *Managing when vendor and supplier risk becomes your own*. McKinsey & Company. https://www.mckinsey.com/capabilities/risk- and-resilience/our-insights/managing-when-vendor-and-supplier-risk-becomes-your-own

Tursunalieva, A., Alexander, D. L. J., Dunne, R., Li, J., Riera, L., & Zhao, Y. (2024). Making sense of machine learning: A review of interpretation techniques and their applications. *Applied Sciences, 14*(2), 496.

Zeng, L., Gupta, R., Motwani, D., Yang, D., & Zhang, Y. (2025). Worse than zero-shot? A fact-checking dataset for evaluating the robustness of RAG against misleading retrievals. arXiv. https://arxiv.org/abs/2502.16101