Proceedings of the 11th North American International Conference on Industrial Engineering and Operations Management, Milwaukee, Wisconsin, USA, June 9-11, 2026

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/NA11.20260117
Published: June 9, 2026

Beyond the Horizon: How Digital Twins and Agentic AI Are Reinventing Enterprise Supply Chain Management

Jacob Davison, Kangmin Liu, Xiaoyang Liu and Yuda Shao
Data Scientist
Ford Global Data Insight & Analytics
Ford Motor Company
Dearborn, MI, USA
jdavi720@ford.com, kliu69@ford.com, xliu272@ford.com, yshao22@ford.com

Chaoye Pan
Technical Expert – Supply Chain Risk
Ford Global Data Insight & Analytics
Ford Motor Company
Dearborn, MI, USA
cpan9@ford.com

Joseph Wyrzykowski
Supply Operations Senior Manager
Ford NA Consumer Business Group Purchasing
Ford Motor Company
Dearborn, MI, USA
jwyrzyk1@ford.com

Abstract

Digital twins and agentic Artificial Intelligence (AI) are reshaping enterprise supply chain analytics by enabling real-time monitoring, scenario evaluation, and context-aware decision support. In automotive supply chains, where disruptions emerge across complex supplier, logistics, and operations networks, traditional analytics workflows often require users to manually synthesize dashboard metrics, case records, and business documents before action can be taken. SupplyAssistAI addresses this challenge as an LLM-powered intelligent assistant for delivery risk analytics and operational incident management. Given a number of data domains which align with the industry standard Supply Chain Operations Reference (SCOR) model managed by the Association for Supply Chain Management (ASCM), the system combines structured supplier performance signals, including SCOR-aligned supplier performance signals spanning reliability, agility, cost, inventory exposure, and risk indicators, with unstructured enterprise content corresponding to signals modeled by the SCOR Enable process. Its architecture uses a distributed Model Context Protocol (MCP) toolset and an orchestration layer to modularize business logic, retrieve relevant data, and generate actionable insights through natural language interaction. SupplyAssistAI also supports multi-level reporting for analyst, management, and executive audiences, enabling a consistent flow from detailed operational analysis to higher-level summaries. Beyond the specific use case, the system serves as a proof of concept for a scalable, decentralized, multi-agent enterprise AI architecture in supply chain analytics.

Keywords
Agentic AI, Supply Chain Analytics, Delivery Risk Analytics, Model Context Protocol, Multi-Agent Systems

1. Introduction
Automotive supply chains are complex systems composed of suppliers, transportation networks, production operations, inventory constraints, and business processes that must remain synchronized under volatile conditions. When supplier disruptions occur, business users often need to navigate multiple dashboards, incident resolution artifacts, performance metrics, and policy documents, before they can determine the current state of risk and identify an appropriate response. This fragmented workflow slows decision-making and increases the burden on both analytics teams and business stakeholders.

Recent developments in digital twins (Ivanov and Dolgui 2021; Ivanov 2023; Ivanov 2024; Ivanov and Gusikhin 2026; van der Valk et al. 2022; Le and Fan 2024; Zheng et al. 2022; Badakhshan and Ball 2023) and Large Language Model (LLM)-based agentic systems (Yao et al. 2023; Xu et al. 2023; Yang et al. 2023; Jackson et al. 2024) offer a new way to address this problem. A digital twin provides a real-time, data-driven representation of the operational environment, while agentic AI enables reasoning, retrieval, summarization, and workflow support on top of that environment. In this work, SupplyAssistAI is presented as an agentic orchestration interface that connects to and coordinates with digital twin data layers. Together, these capabilities are designed to transform analytics from passive dashboard consumption into interactive, context-aware decision support.

SupplyAssistAI is designed as an LLM-powered intelligent assistant for supply chain operations. Its purpose is to help users navigate supplier disruption cases from open to close by integrating structured supplier performance data and unstructured enterprise knowledge into a unified conversational interface. The system also serves as a proof of concept for a scalable, distributed, multi-agent architecture that can support broader supply chain analytics across the enterprise ecosystem.

1.1 Objectives
The objectives of this study are:

- To define the business purpose and operating logic of SupplyAssistAI.
- To describe the system architecture, including the roles of MCP servers and the orchestration layer.
- To explain how SupplyAssistAI combines structured and unstructured data for delivery risk analytics.
- To present the major features of the system, including case management support, reporting, and extensibility to custom and scenario-planning assistants.
- To position SupplyAssistAI as an actively deployed realization of a scalable pattern for decentralized enterprise AI in supply chain analytics.

2. Literature Review
The development of SupplyAssistAI builds on three key areas: supply chain risk analytics, supply chain digital twins, and agentic AI systems. Together, these domains enable the transition from descriptive analytics toward integrated, decision-oriented enterprise systems.

Supply chain risk analytics has established the importance of understanding disruption impact across complex, multi-tier networks. Simchi-Levi et al. (2015) developed a risk-exposure modeling framework that evaluates the operational impact of disruptions without requiring prior probability estimation. This approach enables firms to prioritize mitigation strategies based on impact, improving resource allocation and resilience planning. More recent work has explored predictive modeling approaches for supply chain disruption using heterogeneous time series data (Do et al. 2024). However, while such models provide valuable outputs, their integration into real-time decision workflows remains limited and constrained by the ability to translate into actionable business insights.

Supply chain digital twins (SCDTs) provide a structural framework for operationalizing these insights. Digital twins are dynamic, data-driven representations of physical systems that integrate real-time data, analytics, and simulation for decision support (Ivanov and Dolgui 2021; Ivanov 2023; Ivanov 2024; Ivanov and Gusikhin 2026). They enable end-to-end visibility, scenario analysis, and resilience stress testing across supply chain networks (Zheng et al. 2022; van der Valk et al. 2022; Le and Fan 2024). Additional research highlights the role of digital twins in disruption detection, financial flow management, and resilience optimization (Badakhshan and Ball 2023; Ashraf et al. 2024; Li et al. 2025). Despite these advances, many implementations remain fragmented, with analytics distributed across disconnected tools and teams.

Advances in artificial intelligence, particularly agent-based systems and large language models, offer a mechanism to address this fragmentation. Agent-based approaches support distributed decision-making across complex systems (Sterling and Taveter 2009; Arunachalam and Sadeh 2005), while recent LLM-based methods enable reasoning, tool use, and interaction with both structured and unstructured data (Yao et al. 2023; Xu et al. 2023; Yang et al. 2023). Recent work also highlights the role of generative AI in democratizing access to optimization and decision-making capabilities in enterprise settings (Simchi-Levi et al. 2025). When combined with digital twins, agentic AI enables systems to move beyond passive analysis toward active orchestration of data, models, and workflows (Ivanov and Gusikhin 2026; Jackson et al. 2024; Chen et al. 2025).

To support enterprise-scale deployment, these capabilities require a decentralized architectural approach. In this framework, domain-specific knowledge is modularized into independent components that encapsulate data access, analytics, and business logic, while an orchestration layer coordinates their interaction. This enables scalable integration of distributed expertise and heterogeneous data sources, aligning with the structure of large organizations.

Despite these advances, a gap remains in translating digital twin and risk analytics capabilities into unified, user-facing systems for supply chain operations. SupplyAssistAI addresses this gap by combining digital twin principles, predictive risk analytics, and a decentralized multi-agent architecture to support delivery risk analytics and operations management. By integrating structured and unstructured data through an orchestration layer, the system enables context-aware, scalable decision support in enterprise supply chains.

3. Methods

This study proposes the design of a scalable framework for an intelligent assistant for delivery risk analytics operations management. Rather than evaluating a single predictive model, this study defines the architecture, logic, and workflow of SupplyAssistAI as an actively deployed, production-integrated system combining digital twin concepts, agentic AI, and decentralized multi-agent design to serve as a functional proof-of-concept.

SupplyAssistAI is built as a **conversational, LLM-powered assistant** that interacts with a distributed set of domain-specific services through a structured orchestration layer. The system architecture consists of three primary components: (1) a user interface, (2) an LLM-based orchestration layer, and (3) a distributed Model Context Protocol (MCP) toolset.

The user interface, implemented as a web-based application, enables business users to interact with the system using natural language queries. These queries are passed to the orchestration layer, which is responsible for interpreting user intent, selecting appropriate tools, and coordinating the response generation process. The orchestration layer leverages function-calling capabilities and agent workflows to dynamically route requests across multiple data sources and analytical modules.

The MCP toolset represents the core of the proposed decentralized architecture. Each MCP server encapsulates domain-specific knowledge, including data retrieval logic, analytical rules, and prompt generation templates. MCP servers are organized by SCOR-aligned functional domains (Association for Supply Chain Management 2025) such as supplier delivery reliability (RL.2.6), supply chain cost (CO.1.1), source agility (AG.1.1/AG.2.2), inventory exposure (AM.2.2), risk exposure (AG.1.4), supplier information, and Enable-domain knowledge retrieval. Each server exposes three primary types of functionalities: (1) structured data retrieval, (2) prompt generation for domain-specific analysis, and (3) LLM-based insight generation.

The system operates through a **query routing and orchestration workflow**. When a user submits a request, the orchestration layer identifies relevant MCP domains, retrieves structured and unstructured data, and combines results into a coherent response. Conversation history is maintained to provide contextual continuity and enable follow-up queries.

This architecture, as depicted in Figure 1, reflects a **decentralized multi-agent framework**, in which analytical logic is distributed across modular components rather than embedded in a single monolithic system. The orchestration layer acts as a coordinating agent, dynamically integrating outputs from specialized domain agents. This approach enables scalability, modularity, and reuse of domain knowledge across applications, while supporting flexible integration of new data sources and analytical capabilities.

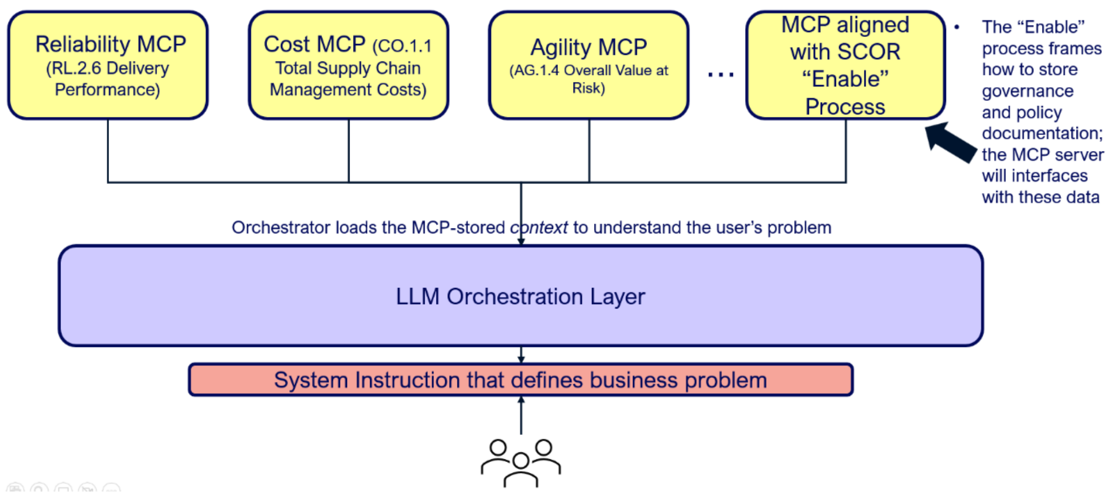

Figure 1. Schematic architecture of SupplyAssistAI. Business logic is captured in MCP servers and dynamically loaded to the orchestrator at completion time.

4. Data Collection

SupplyAssistAI integrates both **structured and unstructured data sources** to support operational incident management. The system is built on top of the supplier operations digital twin, which provides a consolidated view of supplier performance and disruption-related information.

Structured data are derived from key performance domains used in delivery risk assessment. These include SCOR-aligned supplier delivery reliability indicators, supply chain cost indicators, source agility measures, supplier segmentation, inventory exposure signals, and modeled risk exposure indicators. These data sources provide quantitative measures of supplier performance and operational risk and are accessed through domain-specific MCP servers.

In addition to structured data, the system incorporates a wide range of **unstructured enterprise data**, including documentation and records of problem descriptions, and corrective action artifacts such as interim corrective actions (ICA) and permanent corrective actions (PCA). These data are indexed using a search-based retrieval system, enabling semantic access and contextual integration into analytical workflows.

The data architecture follows a **federated model**, in which each MCP server connects to a specific data domain while exposing standardized interfaces for retrieval and analysis. This design enables the system to combine heterogeneous data sources without requiring full centralization, supporting scalability and modular expansion.
Conversation history is also treated as a critical data component. By maintaining prior interactions, the system enables context-aware responses, follow-up analysis, and iterative refinement of insights. This capability is particularly important in case management scenarios, where users may progressively explore different aspects of supplier risk. Table 1 summarizes the data domains, and their analytical purpose, captured by SupplyAssistAI.

Table 1. SupplyAssistAI data domains.

Domain	Data type	Example content	Analytical purpose
Supplier Delivery Reliability (RL.2.6)	Structured	Supplier commit-date adherence indicators, late delivery exception signals	Detect reliability and schedule adherence issues
Supply Chain Cost (CO.1.1)	Structured	Expedite-related cost signals, transportation cost exceptions	Identify cost stress and fulfillment-related cost exposure
Source Agility (AG.2.2)	Structured	Supplier capacity utilization, surge capability indicators	Evaluate supply flexibility and responsiveness to changing demand
Risk Exposure (AG.1.4)	Structured	Model-based disruption signals, leading risk indicators	Anticipate future supplier issues and prioritize proactive intervention
Inventory Exposure (AM.2.2)	Structured	Inventory depletion risk, projected shortage signals	Assess near-term operational exposure and runout vulnerability
Enable Knowledge	Unstructured	Policies, procedures, reference documents, operating guidance	Provide procedural and business context to support analytics and recommendations
Enable Case Resolution	Structured and unstructured	Historical incident records, status summaries, ICA/PCA artifacts	Understand case history, prior actions, and resolution pathways both quantitatively and semantically

In addition to descriptive performance metrics, SupplyAssistAI incorporates outputs from a predictive risk modeling system designed to proactively identify supplier disruption risk. This component represents a key analytical input, enabling forward-looking risk assessment within the business-facing LLM application.

The predictive model is based on an Explainable Artificial Intelligence (XAI) framework that analyzes supplier-level data at a weekly granularity, building on previous work from Do et al. 2024. The modeling pipeline integrates multiple structured data sources, including production releases, capacity metrics, shipment data, behind-schedule indicators, premium freight records, incident resolution records, and historical production loss events. By design, the predictive model is trained in the same structured SCOR-aligned data domains as connected to SupplyAssistAI. These data are processed through an automated machine learning pipeline deployed on a cloud-based infrastructure, where data preprocessing and feature engineering are performed prior to model training and inference.

A key feature of this component is the integration of explainability techniques. SHAP (SHapley Additive exPlanations) values are used to identify the key drivers of each prediction, and these outputs are further translated into natural language explanations using SupplyAssistAI, by combining predictive risk signals with contextual data and case information to support proactive decision-making.

5. Results and Discussion

5.1 System Capabilities and Functional Results

The primary result of this study is the development of SupplyAssistAI as a decentralized **agentic AI system** for delivery risk analytics. The system demonstrates how decentralized multi-agent architecture can be applied to integrate structured performance metrics and unstructured enterprise knowledge into a unified, conversational interface.

SupplyAssistAI supports three core functionalities which are detailed in the following sections.

Supplier Performance and Risk Analysis The system enables users to analyze supplier performance across multiple key performance indicators, including SCOR-aligned supplier performance signals spanning reliability (RL.2.6), agility (AG.1.1/AG.2.2), cost (CO.1.1), inventory exposure (AM.2.2), and risk exposure (AG.1.4). By aggregating data from multiple MCP domains, SupplyAssistAI provides a comprehensive view of supplier health and operational risk. Each data domain in Table 1 is embedded in an MCP server with a curated set of tools designed, with business feedback, to support supply chain operations teams through the case resolution process. Following the framework explained in Section 3, these MCP toolsets provide SupplyAssistAI access to actionable slices of historical performance data (e.g., weekly top reliability exceptions and highest cost exceptions across supplier operations) and risk analysis data (e.g., weekly top most critical indicators of disruption risk). These data slices are also integrated into explicit analysis instructions (broadcast by each MCP server) which guide the LLM how to interpret those data which are presented. The combination of data and interpretation allows the LLM to help users identify potential cross-domain relationships as a conversation proceeds. Figure 2 demonstrates how the system performs a single-pass risk analysis on a single supplier. User input triggers several LLM reasoning steps in order: 1) the user question is routed to the most appropriate data domain to the answer the question, 2) the corresponding data retrieval tools are deployed by the LLM, 3) another reasoning step distills the extracted data insights according to pre-defined system instruction, 4) the final output is sent the user, including follow up suggestions.

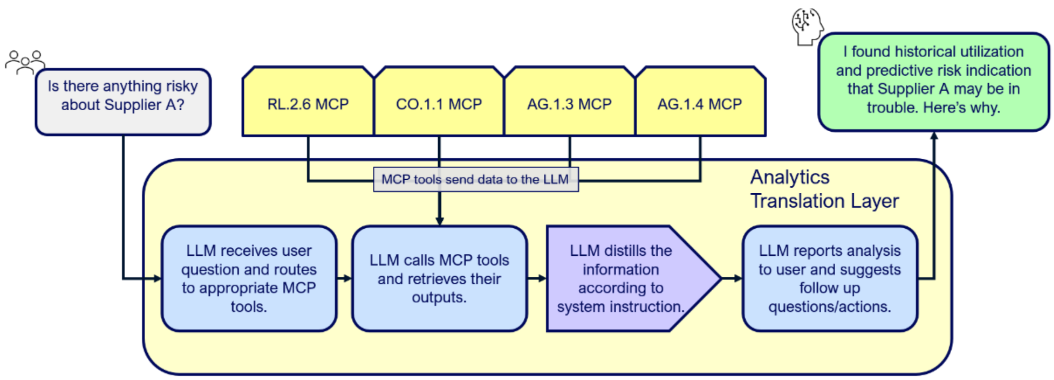

Figure 2. Diagram illustrating an interaction pattern, not proprietary logic, of how a typical analysis question is processed and answered by the system.

Case Management Support The system integrates operational incident data with performance metrics, allowing users to investigate disruption cases from initiation to resolution. Case management data is a combination of structured and unstructured data, and relevant tools which connect the LLM to these data resources are stratified according to their data type. The result is that, like the MCP servers which capture supplier performance and risk data, the case management MCP server broadcasts knowledge of quantitative indicators such as case status and time spent per status, as well as incident records and ICA/PCA artifacts. As the user proceeds through a conversation exploring these data sources, SupplyAssistAI receives context of current open cases, open action items, historical closed cases, historical actions taken, and lessons learned—the combination of incident records with supplier performance and risk analyses allows SupplyAssistAI to provide decision-support suggestions regarding potential next actions that may help move a case toward resolution. Figure 3 demonstrates a typical conversation flow to help the user define the next best action toward moving a case forward. Note that the answer to each question in the session follows an analytics translation procedure similar to Figure 2. In this example, a user asks the system to review the current incident environment for Supplier A. The conversation is designed to proceed toward highlighting concerning data patterns.

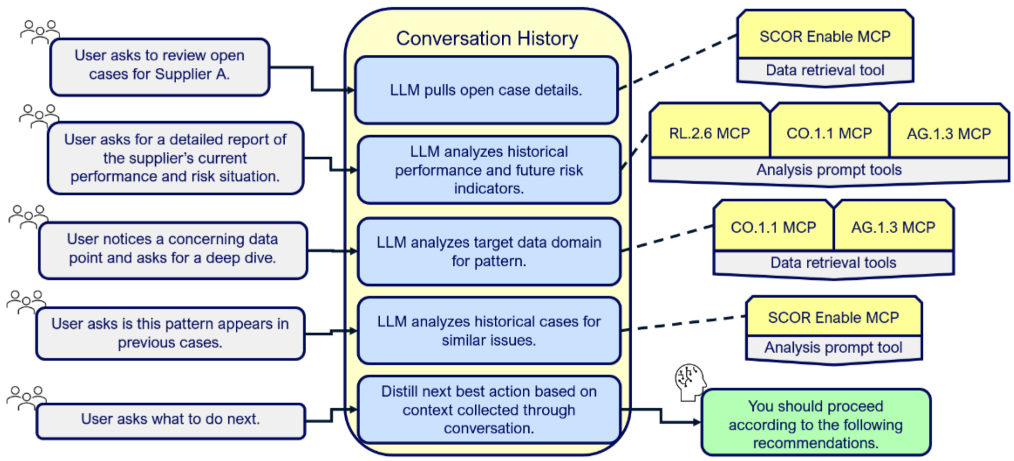

Figure 3. Diagram illustrating an interaction pattern, not proprietary logic, of how a case management conversation proceeds and how the system gathers context based on user queries.

Multi-Level Reporting Analyst teams rely on detailed, near-real-time data to manage day-to-day operations, whereas project management requires performance summaries to monitor processes and improve efficiency. Executive leadership, by contrast, depends on high-level strategic reports to support decision-making and long-term planning. To address these diverse reporting needs, SupplyAssistAI implements a reporting pipeline designed to support role-specific insights across organizational levels, Figure 4 schematically illustrates this reporting pipeline. In the first stage, an analyst submits a request for a detailed report to SupplyAssistAI. The system then orchestrates the request across multiple information domains, including Supplier Information, Risk Exposure (AG.1.4), Supplier Delivery Reliability (RL.2.6), Supply Chain Cost (CO.1.1), Source Agility (AG.1.1/AG.2.2), Inventory Exposure (AM.2.2), and Enable-domain case management data, to generate the required insights. The resulting analyst report provides daily detailed summaries and tabular views of supplier performance, risk indicators, and case-specific context. Analysts can use this report to identify emerging issues, investigate root causes, and initiate appropriate corrective actions. After the analyst report is generated, management- and executive-level reports are derived from it. We do not regenerate these reports independently from each MCP server for two reasons. First, the analyst report already consolidates comprehensive insights from all relevant MCP domains; generating separate reports from each MCP server would introduce duplication, increase the risk of inconsistencies, and add unnecessary overhead. Second, by generating and pre-storing analyst reports for high-risk suppliers, the system reduces the time and computational cost required to produce higher-level reports, thereby improving responsiveness for management and executive users. The SupplyAssistAI reporting pipeline is designed to reduce the amount of manual navigation required across complex data tables and to help users focus on higher-value interpretation. By embedding domain guidance into LLM prompts and grounding report generation in MCP-accessible data domains, the system aims to produce context-aware summaries that are consistent with available structured and unstructured information. This reporting workflow illustrates how an agentic architecture can support more responsive and standardized analytical communication across organizational levels.

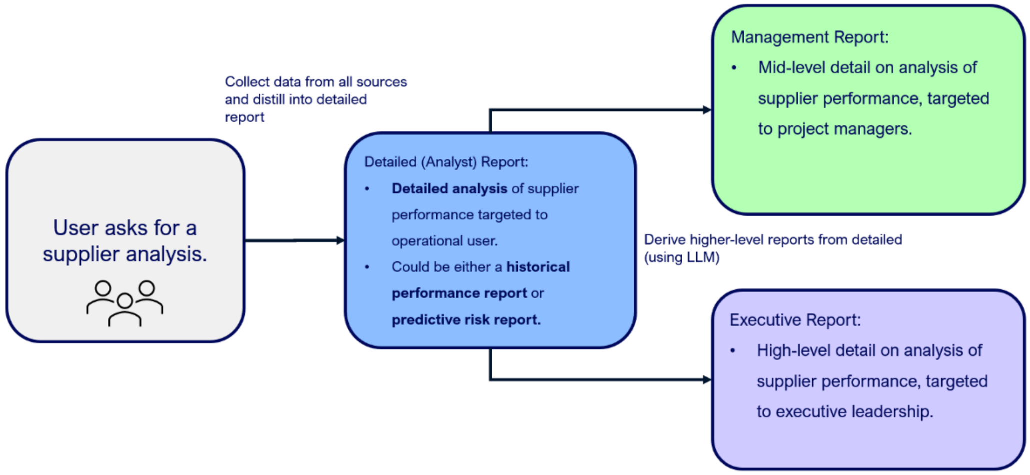

Figure 4. Reporting pipeline showing generation of analyst-level reports and derivation of management and executive summaries.

These capabilities demonstrate a shift from static dashboards to **interactive, context-aware analytics**, where users can dynamically explore supplier issues through natural language interaction.

Moreover, these core capabilities form a platform upon which additional features of SupplyAssistAI may be scaffolded, as illustrated in Figure 5. In addition to the “default” case management assistant, the platform offers additional features including a “Custom Assistant” and a Scenario Planning Assistant.

Custom Assistant This feature is a no-code business solution identity that allows business users to combine process-specific data, via multimodal upload, with the foundational data analysis provided through the MCP data domains and default SupplyAssistAI persona. This allows business users to customize LLM solutions, grounded in the capabilities of SupplyAssistAI, to specific business processes. For example, imagine a business user requires a particular cross-domain analysis of data to be performed every morning on a given supplier. Utilizing this feature, the user may create a Custom Assistant encoding this instruction to a persistent persona and attaching (by upload) a template of how to format the report. This conditions the Custom Assistant perform the same task every time the assistant is called. We note the Custom Assistant may be centrally managed at the business office level and tied to specific business processes that require consistent task automation to support a variety of analyst tasks.

Scenario Planning Assistant This feature connects the foundational base layer to a suite of simulation capabilities that, combined with the MCP data domains, allow business users to analyze “what-if?” scenarios, planning the appropriate mitigation strategy to potential disruption outcomes.

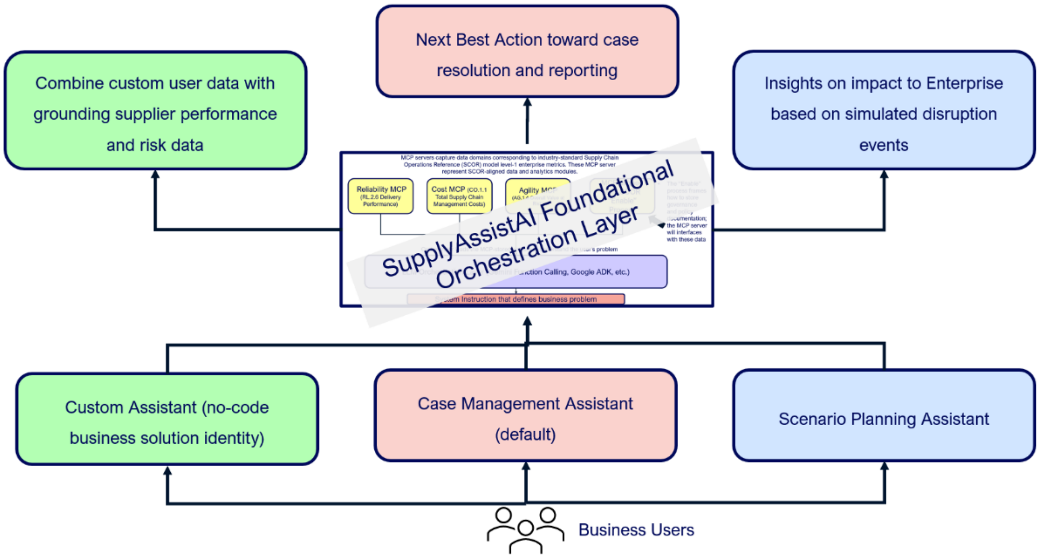

Figure 5. Extension of the core system to additional capabilities, including Custom Assistant and Scenario Planning Assistant built on the foundational analytics layer.

To illustrate how the system operates conceptually, consider a generic, synthetic use-case scenario. A business analyst queries the system: *'What is the risk status of our critical steering column components for next week's production?'* The LLM orchestration layer parses this natural language intent and dynamically routes the request. It first calls the Risk Exposure (AG.1.4) MCP server to retrieve predicted disruption signals for the generic supplier (e.g., 'Supplier A'). Simultaneously, it queries the Inventory Exposure (AM.2.2) MCP server to check current safety stock levels and projected shortages. The system then synthesizes these structured risk signals with unstructured procedural guidelines from the Enable Knowledge database. Finally, SupplyAssistAI generates a unified, context-aware natural language response explaining that Supplier A faces a moderate disruption risk due to sub-tier material constraints, but current safety stock is sufficient to cover next week's production, and recommends initiating an interim corrective action (ICA) tracking case as a proactive measure.

5.2 Architectural and Operational Insights

From an architectural perspective, the system illustrates the practical applicability of a **decentralized, MCP-based design**. By distributing analytical logic across domain-specific servers, each of which capture a data domain in the digital twin, the system enables independent development and maintenance of business logic while supporting integration through a common orchestration layer.

This approach provides several key advantages:

- **Modularity:** New data domains and analytical capabilities can be added by introducing new MCP servers without modifying the core system.
- **Scalability:** The architecture supports expansion across additional supply chain domains and use cases.
- **Reusability:** Domain-specific logic can be reused across multiple applications, enabling knowledge sharing across the enterprise.
- **Interoperability:** Standardized interfaces enable integration with other AI-driven applications and enterprise systems.

The concept of MCP servers forming a “knowledge broadcast layer” highlights the potential for cross-application collaboration, where insights generated in one system can be leveraged by others. This represents a shift toward **enterprise-wide knowledge sharing and distributed intelligence**.

Furthermore, implementing agentic AI in enterprise settings requires robust governance and safety frameworks. SupplyAssistAI is designed to operate within a secure enterprise cloud boundary, utilizing role-based access controls (RBAC) to ensure users only access data domains aligned with their operational clearance. Additionally, the system maintains a 'human-in-the-loop' paradigm, where the assistant provides decision support and recommendations, but all final operational actions and case resolutions require human validation, mitigating risks associated with potential model hallucinations.

5.3 Business Impact and Practical Implications

SupplyAssistAI provides several practical benefits for supply chain operations:

- **Improved accessibility:** The LLM is designed to translate complex analytical and machine-learning model results into natural language business insights, grounded in appropriate knowledge of business policies and procedures.
- **Support for faster decision making:** Single interface for accessing important and timely analysis information aims to facilitate the rapid distribution of knowledge as well as action on critical issues.
- **Potential for emergent analytical insights:** Combining several separate data domains into a single LLM-powered analysis framework helps users identify potential cross-domain patterns across traditional information silos.

6. Conclusion

This paper presented SupplyAssistAI as an agentic AI system for supporting operational incident management in automotive supply chains. The system integrates digital twin concepts, risk analytics frameworks, and a decentralized multi-agent architecture to provide context-aware, conversational decision support.

The proposed architecture distributes domain-specific knowledge across modular MCP servers and uses an orchestration layer to dynamically integrate structured and unstructured data. This approach enables scalable, flexible, and reusable analytics capabilities while aligning with the distributed nature of enterprise systems.
SupplyAssistAI demonstrates how agentic AI can transform supply chain analytics from static dashboards into interactive, intelligent systems that support end-to-end case management. By combining real-time data access, domain-specific reasoning, and multi-level reporting, the system provides a practical pathway toward more responsive and integrated decision-making in supply chain operations.

References

Association for Supply Chain Management, Supply Chain Operations Reference Model: SCOR Digital Standard, ASCM, Chicago, IL, 2025.

Arunachalam, R. and Sadeh, N. M., The supply chain trading agent competition, *Electronic Commerce Research and Applications*, vol. 4, no. 1, pp. 66-84, 2005.

Ashraf, M., Eltawil, A. and Ali, I., Disruption detection for a cognitive digital supply chain twin using hybrid deep learning, *Operations Research International Journal*, vol. 24, no. 23, 2024.

Badakhshan, E. and Ball, P., Applying digital twins for inventory and cash management in supply chains under physical and financial disruptions, *International Journal of Production Research*, vol. 61, no. 15, pp. 5094-5116, 2023.

Chen, W., Choi, T. M., Dolgui, A., Ivanov, D. and Pesch, E., Digital manufacturing and supply chain: creating benefits through operations research and artificial intelligence, *Annals of Operations Research*, vol. 344, pp. 567-574, 2025.

Do, B., Li, X., Pan, C. and Gusikhin, O., Forecasting automotive supply chain disruption with heterogeneous time series, *arXiv preprint arXiv:2407.16739*, 2024.

Ivanov, D., Conceptualisation of a 7-element digital twin framework in supply chain and operations management, *International Journal of Production Research*, vol. 62, no. 6, pp. 2220-2232, 2024.

Ivanov, D., Intelligent digital twin (iDT) for supply chain stress-testing, resilience, and viability, *International Journal of Production Economics*, vol. 263, pp. 108938, 2023.

Ivanov, D. and Gusikhin, O., Supply chain digital twin design and implementation at scale: A case study at the Ford Motor Company and generalizations, *Omega*, vol. 139, pp. 103447, 2026.

Ivanov, D. and Dolgui, A., A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0, *Production Planning & Control*, vol. 32, no. 9, pp. 775-788, 2021.

Jackson, I., Ivanov, D., Dolgui, A. and Namdar, J., Generative artificial intelligence in supply chain and operations management: a capability-based framework for analysis and implementation, *International Journal of Production Research*, vol. 62, no. 17, pp. 6120-6145, 2024.

Jahani, H., Jain, R. and Ivanov, D., Data science and big data analytics: a systematic review of methodologies used in the supply chain and logistics research, *Annals of Operations Research*, 2023.
Le, T. V. and Fan, R., Digital twins for logistics and supply chain systems: literature review, conceptual framework, research potential, and practical challenges, *Computers & Industrial Engineering*, vol. 187, pp. 109768, 2024.
Li, X., Krivtsov, V., Pan, C., Nassehi, A., Gao, R. X. and Ivanov, D., End-to-end supply chain resilience management using deep learning, survival analysis, and explainable artificial intelligence, *International Journal of Production Research*, vol. 63, no. 3, pp. 1174-1202, 2025.
Simchi-Levi, D., Schmidt, W., Wei, Y., Zhang, P. Y., Combs, K., Ge, Y., Sanders, M. and Zhang, D., Identifying risks and mitigating disruptions in the automotive supply chain, *Interfaces*, vol. 45, no. 5, pp. 375-390, 2015.
Simchi-Levi, D., Dai, T., Menache, I. and Wu, M., Democratizing optimization with generative AI, *SSRN Working Paper*, 2025.
Sterling, L. S. and Taveter, K., *The Art of Agent-Oriented Modeling*, MIT Press, Cambridge, MA, 2009.
Van der Valk, H., Strobel, G., Winkelmann, S., Hunker, J. and Tomczyk, M., Supply chains in the era of digital twins – a review, *Procedia Computer Science*, vol. 204, pp. 156-163, 2022.
Xu, B., Peng, Z., Lei, B., Mukherjee, S., Liu, Y. and Xu, D., ReWOO: Decoupling reasoning from observations for efficient augmented language models, *arXiv preprint arXiv:2305.18323*, 2023.
Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K. and Cao, Y., ReAct: Synergizing reasoning and acting in language models, *arXiv preprint arXiv:2210.03629*, 2023.
Yang, H., Yue, S. and He, Y., Auto-GPT for online decision making: benchmarks and additional opinions, *arXiv preprint arXiv:2306.02224*, 2023.
Zheng, X., Lu, J. and Kiritsis, D., The emergence of cognitive digital twin: vision, challenges and opportunities, *International Journal of Production Research*, vol. 60, no. 24, pp. 7610-7633, 2022.

Biographies

Dr. Jacob Davison graduated from Central Michigan University in 2018 with a degree in Physics and Astronomy. He completed his doctoral training in computational Nuclear Theory at Michigan State University, where he earned a dual PhD in Physics and Computational Science in 2023. He joined Ford Motor Company as a research data scientist shortly after, where he leverages his background in AI/ML to solve complex analytical problems, and design generative AI solutions, in supply chain operations.

Dr. Kangmin Liu is a Data Scientist at Ford Motor Company, specializing in supply chain analytics, predictive modeling, and machine learning applications. He holds a Ph.D. in Computational Chemistry from Wayne State University and has over five years of experience applying data-driven methods to large-scale industrial systems. His work spans connected vehicle analytics and supply chain operations, supporting data-informed decision-making and operational efficiency.

Dr. Xiaoyang Liu is a Data Scientist at Ford Motor Company with over 13 years of experience driving transformative innovation in the automotive industry. Holding a Ph.D. in Computational Physics from the University of Tokyo, alongside Master's and Bachelor's degrees in Engineering Physics from Tsinghua University, Dr. Liu specializes in architecting end-to-end AI/ML solutions and Agentic workflows across supply chain, quality, reliability, and ADAS. At Ford, Dr. Liu engineered a comprehensive AI/ML framework for supply chain delivery risk modeling alongside an AI-driven simulation system, leveraging predictive modeling and Explainable AI (XAI) to distill hundreds of complex variables into actionable KPIs. Previously at General Motors, he earned the "Critical Technical Talent Award" for optimizing ADAS sensor alignment for SuperCruise and developing highly efficient, automated algorithm benchmarking platforms. His tenure at Stellantis (FCA) was highlighted by the creation of a Warranty Early Warning System that drove massive cost reductions, as well as the optimization of pickup reliability testing plans using connected vehicle big data. A prolific inventor with four patents filed in camera systems and AI applications, Dr. Liu seamlessly bridges complex physics-based modeling with executive strategy. He leverages deep expertise in cloud platforms, machine learning, computer vision, and deep learning. Backed by a DFSS Black Belt, multiple excellence awards, and over 220 citations for pioneering research, Dr. Liu consistently translates advanced analytics into high-impact business strategies.

Dr. Yuda Shao graduated from The Ohio State University in 2020 with a degree in Actuarial Science before pursuing doctoral studies. In 2025, he earned his Ph.D. in Statistics from the University of Virginia and subsequently joined Ford as a Data Scientist, where he applies his advanced analytical expertise.

Dr. Chaoye Pan is a Data science & Machine Learning Manager at Ford Global Data Insight & Analytics, where he leads supply chain risk analytics team. He has over 10 years of experience in application of advanced technology and analytics in the automotive industry. During his service at Ford, he has created numerous high-impact long-lasting applications for Ford manufacturing and supply chain focusing on Gen AI, machine learning, deep learning, and optimization areas.

Joseph Wryzykowski is a senior leader in supply chain operations and digital intelligence at Ford Motor Company, specializing in AI-driven transformation, digital twins, and operational resilience. Joseph bridges supply chain operations and digital product management, leading a distributed, global, cross-functional organization to deliver risk sensing, predictive intelligence, and decision systems that protect production continuity and reduce systemic risk. Joseph owns the enablement of Ford's Supply Chain Graph and Digital Twin, embedding advanced analytics and delivery-risk models into standard operating cadence. He also leads a distributed global team responsible for deep-tier value-stream mapping, powering impact assessment and real-time risk sensing across the supply network. Through deep-tier discovery, scenario analysis, and mitigation planning, Joseph has pioneered the application of agentic AI and decentralized analytics architectures to enable real-time disruption detection, faster executive decision-making, and proactive response across complex supplier ecosystems. Recognized for strategic leadership and cross-functional collaboration, his work has accelerated adoption of resilient supply chain practices and enterprise digital capabilities. Joseph holds advanced degrees in engineering and business and is committed to advancing industry standards through innovation, operational excellence, and knowledge sharing.