International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 6, Issue 4, 49-58, 2025
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V6I4P107

Original Article

Architecting, Grounding and Governing – Telecom Supply Chain GenAI Agents with SAP AI Foundation

Sivasubramanian Kalaiselvan
Distinguished Engineer, USA.

Received On: 28/08/2025 Revised On: 02/10/2025 Accepted On: 09/10/2025 Published On: 28/10/2025

Abstract - *Modern supply chains demand unprecedented agility, resilience, and intelligence to navigate the volatility of rapid technology cycles, complex network rollouts, and critical infrastructure dependencies. Generative AI (GenAI) agents offer transformative potential, automating complex decision-making and optimizing processes. However, realizing this potential requires robust frameworks for their architecture, grounding in reliable data, and responsible governance. This white paper details a comprehensive approach to architecting, grounding, and governing Telecom Supply chain GenAI agents leveraging SAP AI Foundation on the SAP Business Technology Platform (BTP). It explores methodologies for building context-aware, reliable, and ethically aligned agents using components like SAP AI Core, Generative AI Hub, SAP HANA Cloud Vector Engine, and SAP Knowledge Graph. The paper outlines the benefits, including enhanced efficiency, enhanced network resilience, improved field service efficiency, optimized inventory, improved decision-making, and increased resilience, alongside implementation considerations and use cases, targeting supply chain professionals, IT/AI architects, and business leaders seeking to harness GenAI responsibly within the SAP ecosystem and architecting Gen AI agents.*

Keywords - *Telecom Supply Chain Management (SCM), Generative AI (GenAI), AI Agents, SAP AI Foundation, SAP Business Technology Platform (BTP), Architecting, Grounding, Governing, SAP AI Core, Generative AI Hub, Joule Studio, SAP HANA Cloud Vector Engine, SAP Knowledge Graph, Retrieval Augmented Generation (RAG), AI Ethics, Responsible AI, Supply Chain Optimization, Digital Transformation, Network Rollout, Spare Parts Logistics, 5G.*

1. Introduction

The contemporary telecom supply chain operates within a dynamic environment characterized by global disruptions, shifting consumer demands, and intricate partner networks. Traditional management approaches struggle to cope with this complexity, often resulting in inefficiencies like the bullwhip effect, suboptimal inventory levels, and delayed responses to market changes [1]. Enterprise Resource Planning (ERP) systems, while central, face challenges in integrating diverse data and enabling proactive decision-making. The advent of Generative AI (GenAI), particularly large language models (LLMs) and agentic AI systems, presents a paradigm shift. Agentic AI refers to systems capable of autonomous decision-making and action towards specific goals with limited human intervention. These agents can sense their environment, reason through complex scenarios, learn from outcomes, and interact with systems and humans, offering a pathway to self-optimizing supply chains.

SAP, recognizing this potential, provides SAP AI Foundation on the SAP Business Technology Platform (BTP) as a comprehensive suite of tools and services to build, manage, and govern AI solutions, including GenAI agents. SAP's strategy emphasizes AI that is Relevant, Reliable, and Responsible, ensuring that AI adoption drives business value while adhering to ethical standards and grounding in enterprise context. This white paper outlines how organizations can leverage SAP AI Foundation to effectively architect, ground, and govern GenAI agents for transformative impact in telecom supply chain management.

2. Problem Statement

Supply chain management faces persistent challenges that hinder efficiency, resilience, and profitability. Key issues include Network Rollout and Maintenance Complexity managing the build-out of 5G/6G networks involves coordinating thousands of components (antennas, basebands, fiber) to thousands of sites, each with unique configurations. Simultaneously, maintaining existing networks requires a high-stakes spare parts logistics operation where a single failure (a "network-down" event) can cost millions per hour. Globalized networks, intricate dependencies, and frequent disruptions (geopolitical events, climate change, pandemics) make forecasting and planning exceedingly difficult. Critical data often resides in disparate systems (ERP, WMS, TMS, supplier portals), Network Management Systems (NMS), Field Service Management (FSM) tools, supplier portals, and unstructured maintenance logs preventing a holistic view and hindering real-time, data-driven decisions. Traditional methods struggle to integrate and leverage the sheer volume and variety of data [2]. The supply chain must manage two distinct demand patterns: highly volatile, promotion-driven demand for customer premises equipment (CPE) like handsets and routers, and complex, project-based demand for network infrastructure. Both are subject to rapid technological obsolescence. Achieving consensus and coordination among multiple stakeholders (internal departments, suppliers, logistics providers) is time-consuming and prone to failures, leading to suboptimal outcomes like the bullwhip effect [1]. Valuable insights are often locked in unstructured formats (emails, reports, news feeds, sensor data), which traditional analytical tools cannot easily process. Deploying autonomous systems requires strong governance frameworks to ensure reliability, security, compliance, and ethical operation. Lack of trust hinders adoption [3] GenAI agents offer a potential solution, but their effective deployment necessitates addressing how to build them reliably (architecture), ensure their outputs are factually correct and contextually relevant (grounding), and manage their operation responsibly (governance).

.

3. Capabilities and Literature Review

The development of effective supply chain GenAI agents hinges on leveraging advanced AI capabilities within a structured framework.

Agentic AI Capabilities include operating independently within defined guardrails to achieve goals, pursuing specific objectives (e.g., minimize network downtime, reduce inventory holding costs for spare parts, ensure 95% on-time delivery for new site build), perceiving changes through data inputs (IoT, ERP, external feeds), analyzing information, decomposing tasks, and planning multi-step actions using foundation models and other techniques, improving performance over time based on feedback and new data, adapting to changing conditions, utilizing APIs, calling external tools (e.g., optimization solvers, databases), and collaborating with humans or other agents.

3.1. SAP AI Foundation Capabilities:

SAP AI Foundation provides the technical backbone:

- SAP AI Core: A service on SAP BTP for training, deploying, and managing AI models (including custom and partner models) throughout their lifecycle. It handles orchestration, scalability, and monitoring.
- Generative AI Hub: Provides secure, managed access to a curated library of large language models (LLMs) from SAP and partners (e.g., OpenAI, Anthropic, Google Cloud). It includes tools for prompt engineering, orchestration (grounding, content filters, data masking), and trust/control features [5].
- Joule: SAP's AI copilot, acting as a natural language interface and orchestrator across SAP applications. Joule Agents can coordinate tasks across multiple specialized agents (e.g., finance, supply chain, HR) to execute end-to-end processes.
- Joule Studio (in SAP Build): A low-code/no-code environment for building custom Joule Skills and AI Agents, enabling developers and business users to extend AI capabilities tailored to specific processes [5].

3.2. Grounding Capabilities:

Ensuring reliability requires grounding agents in factual, enterprise-specific data:

- Retrieval Augmented Generation (RAG): A technique where agents retrieve relevant information from a knowledge base before generating a response. This significantly reduces hallucinations and improves factual accuracy.
- SAP HANA Cloud Vector Engine: Enables efficient storage and querying of high-dimensional vector embeddings derived from enterprise data (structured and unstructured). This is fundamental for implementing RAG, allowing agents to find the most relevant context quickly [5].
- SAP Knowledge Graph: Provides a semantic layer over enterprise data, capturing relationships between entities (e.g., a specific component, its compatible router models, the cell sites it's deployed at, and its authorized suppliers). Agents can query the Knowledge Graph to understand complex contexts and validate LLM-generated insights, enhancing trustworthiness [5].
- SAP Business Data Cloud: Provides unified access to business data across SAP applications, crucial for feeding grounding mechanisms.

3.3. Governance Capabilities:

Responsible deployment necessitates strong governance:

- SAP AI Ethics Policy: Based on UNESCO recommendations and SAP's guiding principles (Human Agency, Addressing Bias, Transparency, Quality, Data Protection), providing a framework for ethical development and deployment [4]. It includes assessment processes to identify red lines and high-risk cases.
- Trust & Control in GenAI Hub: Features like content moderation, data masking, access controls, and logging ensure secure and compliant LLM usage [5].
- Transparency: Mechanisms to understand and document agent decisions, capabilities, and limitations, fostering trust [4]. Human oversight remains critical, especially for high-risk decisions

3.4. Literature Insights:

Research highlights the potential of LLM agents for consensus-seeking and coordination in SCM, potentially overcoming limitations of previous agent-based systems by offering lower entry barriers and near-human reasoning [1].

However, challenges related to reliability, prompt engineering, and the need for human-in-the-loop governance persist [3].

4. Discussions

The integration of GenAI agents into high-stakes telecom supply chain, facilitated by platforms like SAP AI Foundation, sparks several critical discussion points. While grounding techniques significantly improve reliability, LLMs can still hallucinate or make errors. Building trust requires robust validation, transparency in decision-making, clear communication of limitations, and effective human oversight mechanisms. The "garbage in, garbage out" principle applies forcefully. Grounding mechanisms depend on accurate, comprehensive, and accessible enterprise data. Organizations must invest in data governance, quality management, and integration across silos [2]. The SAP Business Data Cloud aims to address this.AI models, including LLMs, can inherit biases from their training data. Optimized spare parts inventory (less capital tied up), automated procurement (lower manual effort), and efficient logistics can result in cost reduction. The primary value. Reducing network downtime by minutes can save millions, justifying the investment immediately and keeping up the resiliency and uptime. A telecom can leverage its agent-driven supply chain as a competitive advantage, offering premium B2B services like "guaranteed network resilience" or "proactive infrastructure assurance" to enterprise customers. SAP's AI Ethics framework mandates bias detection and mitigation, but continuous vigilance and diverse data are essential [4]. Agents interacting with multiple systems and external data sources increase the attack surface. Secure API management, robust access controls, data anonymization/masking (as provided in GenAI Hub), and compliance with data privacy regulations (e.g., GDPR) are non-negotiable [4].

Implementing and running GenAI agents involves costs (compute, platform subscriptions, development, training). While potential benefits like efficiency gains up to 3x productivity [5] and cost reduction are significant, organizations need clear metrics to measure ROI and justify investment. Integrating AI agents requires rethinking processes and roles. Human tasks shift from execution to oversight, exception handling, and strategic decision-making. Upskilling and reskilling the workforce are crucial for successful adoption. Connecting agents seamlessly with existing SAP landscapes (S/4HANA, IBP, Ariba, TM, EWM) and external systems requires careful architecture and robust integration capabilities, often leveraging SAP BTP Integration Suite. Addressing these points proactively is key to unlocking the value of GenAI agents in the telecom supply chain responsibly and sustainably.

5. Detailed Explanation: Architecting, Grounding, and Governing

Leveraging SAP AI Foundation provides a structured approach to building and managing supply chain GenAI agents.

5.1. Architecting Supply Chain Agents:

Designing the blueprint for GenAI agents requires careful consideration of their core intelligence, how they interact, the tools they wield, their memory, and how users engage with them. This architecture must seamlessly integrate within the broader SAP ecosystem.

5.1.1. Core Intelligence (LLM Selection):

The choice of the underlying foundation model is a critical architectural decision. The Generative AI Hub within SAP AI Foundation provides access to a diverse portfolio of LLMs, enabling architects to make informed trade-offs. For instance, automating routine tasks like checking order statuses or classifying field service reports or generating standard shipping documents might be efficiently handled by smaller, faster, and more cost-effective models like Llama 3 or Gemini Flash. Conversely, complex tasks involving multi-step reasoning, scenario simulation for risk mitigation, or intricate negotiation with supplier agents might necessitate the advanced capabilities of larger models such as GPT-4 or Claude 3 Opus, despite their higher computational cost which is explained in Figure 1on he comparison of Capability vs Cost. The Hub facilitates experimentation and selection based on empirical performance against specific telecom SCM benchmarks and budgetary constraints.

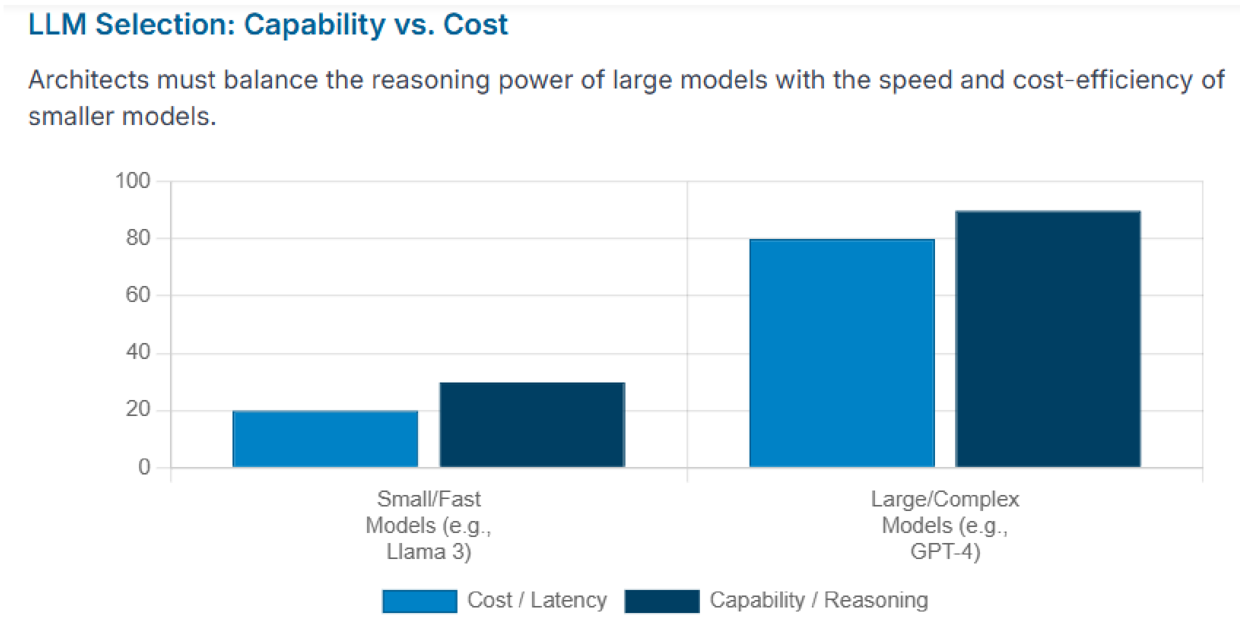

Fig 1: LLM Selection parameter – Capability vs Cost comparison

5.1.2. Agent Framework & Orchestration:

Defining how agents work together is key to automating complex processes. A "monolithic" agent cannot handle the diversity of telecom SCM. The optimal architecture is a multi-agent system, orchestrated by Joule.

Joule Agents: Position Joule as the central orchestrator, the "conductor" for multi-agent workflows. When a user initiates a complex request via Joule (e.g., "Expedite delivery for order X due to customer urgency, considering cost implications"), Joule interprets the intent, identifies the necessary sub-tasks (checking inventory, finding alternative carriers, calculating costs, notifying the customer), and delegates these tasks to specialized agents (Network Rollout Agent, a Logistics Agent, and a Component Risk Agent, Inventory agent, logistics agent, finance agent, communication agent). Figure 2, explains the high-level Joule Agent OS orchestration. Joule manages the communication flow between these agents and synthesizes the results for the user or triggers subsequent actions, leveraging its deep integration across SAP applications [5].

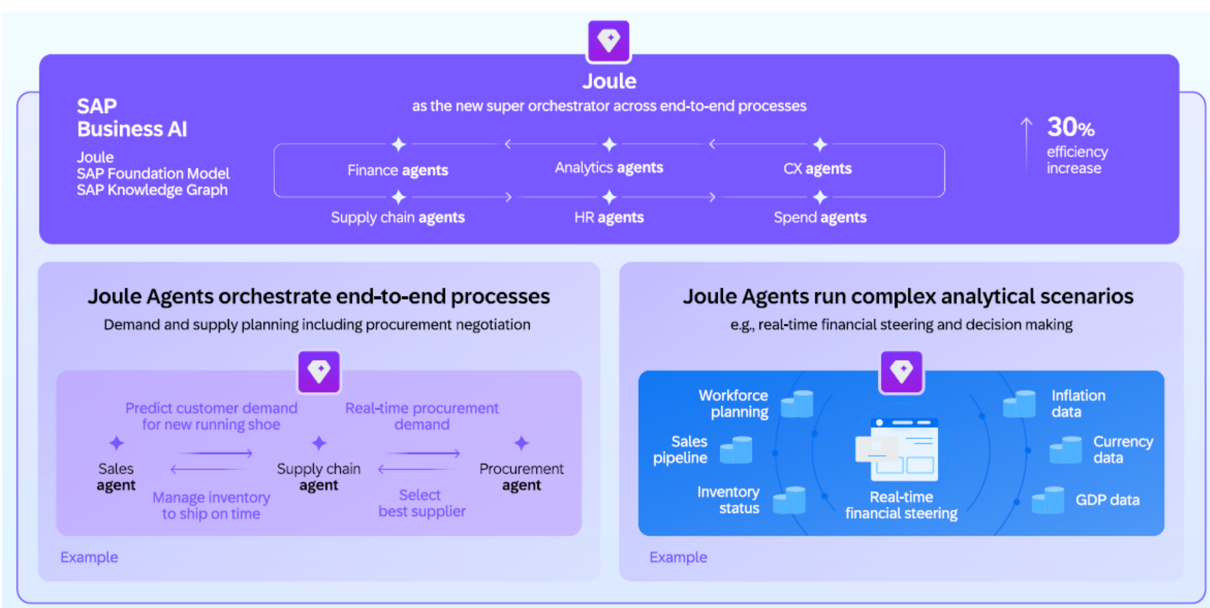

Fig 2: Joule Agent OS Orchestration, courtesy from SAP

Custom Agents: Develop specialized agents using SAP AI Core for fine-tuned control over deployment and operational management. These agents are ideal for encapsulating domain-specific expertise or proprietary algorithms. For example, a custom agent could implement a company's unique spare parts optimization logic for high-value network assets triggered via an API call from Joule or another system. AI Core provides the necessary infrastructure for scalable deployment, versioning, monitoring, and managing the entire lifecycle of these bespoke agents.

Joule Studio: Empower both professional developers and citizen developers using Joule Studio within SAP Build. This low-code/no-code environment democratizes agent creation, allowing business experts (like network planners) to design and build agents or skills that automate specific departmental workflows or address niche requirements [5]. Users can visually define agent triggers (e.g., receipt of a supplier alert), goals (e.g., assess risk and propose mitigation), required tools (e.g., access supplier rating data), and interaction logic, significantly accelerating the development and deployment of tailored AI solutions.

5.1.3. Tools and Connectivity:

Agents derive their power from their ability to interact with the operational environment. Equipping them with the right tools essentially API connections managed securely via SAP BTP Integration Suite, is crucial. This connectivity allows agents to:

- Access Real-time SAP Data: Query SAP S/4HANA for current sales order details, check inventory levels in SAP EWM or IBP, retrieve supplier contract terms from SAP Ariba, or get shipment statuses from SAP Logistics Business Network (LBN) using secure OData APIs or BAPIs. This ensures decisions are based on the latest operational reality.
- Leverage External Services: Integrate external data sources critical for supply chain context, such as real-time weather forecasts impacting transportation, live shipping carrier tracking information, commodity price feeds, or geopolitical risk alerts from third-party providers.
- Invoke Analytical Power: Trigger sophisticated analytical or optimization routines, whether it's using the Predictive Analytics Library (PAL) within SAP HANA Cloud for forecasting, calling external optimization solvers for complex routing problems, or executing custom machine learning models deployed via SAP AI Core.

5.1.4. Memory: Effective agents require memory to maintain context and learn. The architecture should differentiate between:

- Short-Term Memory: Managed within the agent framework (e.g., Joule's conversation context), this allows the agent to track the immediate flow of a conversation or a multi-step task, ensuring coherent interaction without constantly needing to re-establish context.
- Long-Term Memory: Implement persistent storage using robust databases like SAP HANA Cloud. This allows agents to retain knowledge across sessions, such as user preferences, historical performance data of suppliers or carriers, successful resolutions to past exceptions, or learned optimizations. HANA Cloud's capabilities, including its vector engine, can support storing and retrieving complex learned patterns or summarized experiences, enabling agents to improve their performance over time.

5.1.5. User Interface: The way users interact with agents significantly impacts adoption and effectiveness. Seamless integration is key.

- Embedded in Fiori: Integrate agent functionalities directly into relevant SAP Fiori applications. For example, a procurement specialist reviewing purchase requisitions in Fiori might see a proactive suggestion from an agent highlighting a potential sourcing risk, with options to trigger further analysis or mitigation actions directly within the app.
- Conversational via Joule: Utilize Joule's natural language interface as the primary mode of interaction, allowing users to initiate tasks, ask questions, receive summaries, and approve actions across various SAP modules through a unified conversational experience. This lowers the barrier to entry and makes complex functionalities more accessible.

5.2. Grounding Agents in Business Context:

Grounding is the cornerstone of reliable agent performance, ensuring agents provide fact-based, relevant responses deeply connected to the specific enterprise context rather than generic information. This process anchors the agent's knowledge in verifiable company data.

5.2.1. Data Preparation:

The first critical step involves consolidating and preparing high-quality data. This requires breaking down data silos across SAP systems (ERP, SCM, CRM, Ariba) and integrating relevant external feeds (e.g., market reports, logistics updates). Tools like SAP Datasphere facilitate data federation and virtualization, providing a unified view without extensive data movement, while SAP Data Intelligence Cloud enables the creation of robust data pipelines for cleaning, transforming, and harmonizing both structured (e.g., sales orders, inventory levels) and unstructured data (e.g., supplier emails, maintenance reports) into a format suitable for grounding. Accuracy, completeness, and timeliness of this data are paramount.

5.2.2. Vectorization:

Convert the prepared data corpus into vector embeddings using sophisticated models accessed via the Generative AI Hub or custom models deployed on SAP AI Core. These embeddings are numerical representations that capture the semantic meaning of text, documents, images, or other data types. This process transforms diverse information like product specifications, supplier contracts, shipping manifests, historical order patterns, customer feedback, and maintenance logs into a format that enables efficient similarity searches.

5.2.3. Vector Storage & Retrieval:

Store these high-dimensional embeddings efficiently within SAP HANA Cloud's dedicated Vector Engine. When an agent receives a prompt or needs specific context to perform a task, the system converts the query into an embedding and uses it to search the vector store. HANA Cloud rapidly identifies and retrieves the most semantically similar data chunks (e.g., relevant contract clauses, past similar supply chain incidents, specific product details) by comparing vector distances. This ensures the agent receives the most pertinent information almost instantly, crucial for real-time decision support. Retrieval parameters, such as the number of data chunks retrieved, can be tuned for optimal context richness versus conciseness.

5.2.4. RAG Implementation:

Implement the Retrieval Augmented Generation (RAG) workflow, which is orchestrated within the Generative AI Hub. The retrieved data chunks are dynamically inserted into the prompt sent to the selected LLM.
For our technician, the prompt is augmented with:
Context 1: “Eng. Bulletin 72: Part \#47-ABC is end-of-life. It is superseded by Part \#48-XYZ”.

Context 2: “Maintenance Manual v4, Site 1138: WARNING: Must use static-dissipative wrist strap”. This drastically reduces hallucinations and tailors the response. This augmented prompt provides the LLM with specific, factual context directly relevant to the query, drastically reducing the likelihood of generating factually incorrect statements (hallucinations) and ensuring the response is tailored to the company's specific situation. The GenAI Hub manages this complex interaction seamlessly.

5.2.5. Knowledge Graph Integration:

Enhance the grounding further by integrating SAP Knowledge Graph. While RAG retrieves relevant text based on semantic similarity, the Knowledge Graph provides structured understanding of entities and their complex relationships (e.g., supplier A provides component B for product C manufactured at site D, which is impacted by logistics disruption E). Agents can query the graph to understand these intricate dependencies (e.g., "Identify alternative suppliers for component B not impacted by disruption E, considering their lead times and approved status") and to validate information retrieved via RAG or generated by the LLM against established business facts and rules. This adds a layer of semantic validation and deeper contextual awareness.

5.2.6. Prompt Engineering:

Meticulously craft prompts to maximize the effectiveness of grounding. This involves explicitly instructing the LLM to base its reasoning and response *exclusively* on the provided context (retrieved vector data and Knowledge Graph results). Include instructions to cite specific source documents or data points for generated statements, enhancing traceability and user trust.
For example, a prompt might include: "Using *only* the provided context documents$$

Context 1: Supplier Contract Summary, Context 2: Recent Shipment Status$$, identify potential risks for order #12345. Cite the relevant document for each identified risk."

This disciplined approach guides the LLM to function as a context-aware reasoning engine rather than a generic knowledge generator.

Through this multi-faceted grounding process, leveraging SAP's integrated data platform, vector engine, knowledge graph, and AI services, GenAI agents become reliable, trustworthy partners in navigating complex telecom supply chain challenges.

5.3. Governing Agent Operations:

Effective governance is paramount to ensure that GenAI agents operate safely, ethically, compliantly, and maintain user trust throughout their lifecycle. This involves establishing clear rules, implementing technical controls, and fostering a culture of responsibility. Figure 3, explains the ethic assessment process flow.

5.3.1. Ethical Assessment:

The governance journey begins early in the ideation phase with a mandatory ethical assessment using SAP's established AI Ethics Assessment Process. This crucial step proactively identifies potential ethical risks. It involves rigorously checking for alignment with SAP's "red lines" – strictly prohibited uses like discriminatory social scoring or harmful manipulation – and classifying the use case based on potential impact. High-risk scenarios, such as those involving sensitive personal data, fully automated decision-making affecting individuals, or deployment in critical sectors like telecom infrastructure management or healthcare, trigger a mandatory review by the dedicated AI Ethics Steering Committee. This upfront scrutiny ensures that ethical considerations, including fairness, human agency, and potential societal impact, are embedded in the agent's design from the outset. This process is not a one-off check; significant changes in agent scope or functionality necessitate reassessment.

5.3.2. Access Control & Security:

Robust security is fundamental. Utilize the comprehensive security features inherent in the SAP Business Technology Platform (BTP), including strong authentication mechanisms (like SSO) and fine-grained authorization policies. This ensures only legitimate users and systems can interact with the agents and controls precisely what data and functionalities each agent can access within SAP systems (e.g., preventing a field service agent from accessing sensitive financial data). Implementing strict role-based access control (RBAC) is vital for both the agent development environments (like Joule Studio, limiting who can build or modify agents) and the operational runtime (SAP AI Core, controlling deployment and execution permissions). This layered security approach minimizes the risk of unauthorized access, data breaches, or malicious agent manipulation.

5.3.3. GenAI Hub Governance Features:

The Generative AI Hub acts as a central control plane, offering built-in governance tools essential for managing LLM interactions:

- Data Masking/Anonymization: This feature automatically identifies and redacts or anonymizes personally identifiable information (PII) or other sensitive data within prompts before they are transmitted to external LLM providers. This is critical for maintaining data privacy and complying with regulations like GDPR, especially when using third-party models.
- Content Filters: These filters act as safety nets, scanning both user inputs (prompts) and LLM outputs (responses) to detect and block harmful, inappropriate, or policy-violating content (e.g., hate speech, generation of unsafe code) [6]. This helps maintain ethical standards and protect users from potentially damaging outputs.
- Audit Logging: Comprehensive logging captures detailed records of all agent interactions, including user prompts, LLM calls made through the Hub, tools invoked, and final responses. This traceability is essential for monitoring agent behavior, debugging issues, ensuring compliance with internal policies and external regulations, and facilitating post-incident analysis (e.g., reviewing an agent's actions during a network outage).

5.3.4. Human Oversight:

Despite increasing autonomy, maintaining human oversight is a core principle of Responsible AI, especially for decisions with significant consequences. Implement "human-in-the-loop" or "human-on-the-loop" workflows. Design agents to recognize situations requiring human judgment – such as actions involving substantial financial commitments (e.g., approving a large unexpected supplier payment), safety-critical operations (e.g., altering a manufacturing process parameter), or ambiguous scenarios where the agent's confidence score falls below a predefined threshold [7]. The agent should then escalate these situations to designated human operators for review and final approval. SAP's Joule Studio provides tools to visually model and configure these escalation paths and approval workflows, ensuring humans retain ultimate control over critical processes [4].

5.3.5. Monitoring and Performance Management:

Continuous monitoring is vital for ensuring agents perform reliably and effectively over time. Utilize the monitoring capabilities within SAP AI Core and broader SAP BTP services to track key operational metrics such as agent response latency, resource consumption (CPU, memory, token usage), error rates, and API call successes/failures. Beyond operational health, continuously evaluate the quality and accuracy of agent outputs against defined business KPIs (e.g., first-time-fix rate improvements, reduction in network downtime). Implement mechanisms to detect model drift (degradation in performance over time) and monitor for emergent biases in agent decisions. Actively solicit and analyze user feedback through integrated mechanisms to identify areas for improvement and ensure the agent continues to meet business needs effectively. This forms a crucial feedback loop for iterative refinement.

5.3.6. Transparency:

Building and maintaining trust requires transparency. Meticulously document the agent's intended purpose, core capabilities, known limitations, and the specific LLMs it utilizes. Clearly outline the data sources used for grounding the agent's knowledge base. Where technically feasible and appropriate for the use case, provide explanations for significant agent decisions or recommendations [8]. This might involve surfacing simplified LLM reasoning traces (Chain-of-Thought) or highlighting the key data points retrieved via RAG that influenced a particular outcome. For the field technician, the agent citing "Source: Eng. Bulletin 72" is a critical transparency feature that builds immediate trust and allows for verification. Clear documentation mechanisms help users understand how the agent works, build confidence in its outputs, and facilitate troubleshooting.

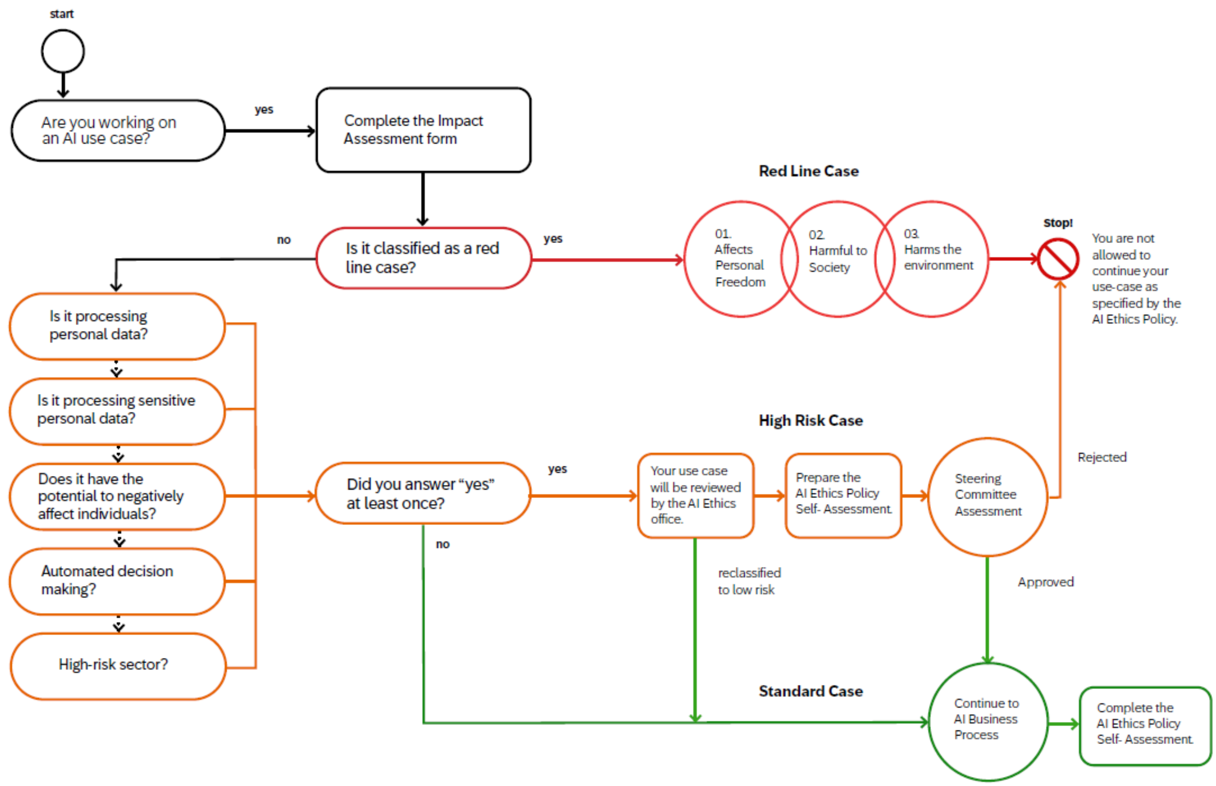

Fig 3: AI Ethics Assessment process, courtesy from SAP

By systematically embedding these governance principles and leveraging the tools within SAP AI Foundation, organizations can deploy powerful supply chain agents that are not only effective but also operate safely, ethically, and in alignment with business values and regulatory requirements. This comprehensive approach transforms governance from a compliance hurdle into an enabler of trustworthy and sustainable AI adoption.

6. Use Cases and Benefits

GenAI agents powered by SAP AI Foundation can revolutionize various telecom SCM functions:

6.1. Intelligent Demand Forecasting:

Agents analyze historical sales (from SAP CAR), DIFs (promotions, events, weather), market trends, and real-time signals (social media, news) to generate more accurate and granular forecasts [2]. They can explain forecast drivers and simulate "what-if" scenarios.

- Benefit: Reduced stockouts/overstocks, improved promotional ROI, optimized inventory

6.2. Autonomous Spare Parts Logistics (Field Service Agent):

As detailed in the grounding example, the agent provides real-time, grounded technical support and automates the logistics of spare parts for field technicians. It predicts part failures based on IoT data from network hardware and maintenance history, pre-emptively shipping parts to forward stocking locations to meet technicians.

- Benefit: Drastically reduced network downtime, improved first-time-fix rates (the technician has the correct, compatible part the first time), and optimized high-value spare parts inventory (reduced holding costs and obsolescence).

6.3. Proactive Network Rollout Management (Network Rollout Agent):

An agent continuously monitors all data related to a 5G network rollout. It senses a supplier's shipment of antennas is delayed by 3 days (from a logistics API). It queries the Knowledge Graph to see which 10 cell site builds are affected. It then simulates the impact in SAP IBP and autonomously re-prioritizes the delivery schedule, moving up 10 other sites that have all their components ready.

- Benefit: Increased agility and resilience. The rollout schedule is dynamically self-optimizing, maximizing build-rate and minimizing costly crew idle time.

6.4. Proactive Risk Management & Response:

Agents continuously monitor global events, supplier news, logistics updates, and IoT sensor data for potential disruptions. They assess impact using the Knowledge Graph, simulate mitigation options (e.g., alternative sourcing, rerouting), and execute responses (e.g., creating stock transfer orders, notifying stakeholders) via Joule orchestration.

- Benefit: Increased supply chain resilience, reduced disruption impact, faster recovery times.

6.5. Intelligent Handset Demand Forecasting (Handset Demand Agent):

Agents analyze historical sales (from SAP CAR), marketing promotions, competitor launch events, and real-time social media sentiment to generate more accurate, granular forecasts for new device launches [2].

- Benefit: Reduced stockouts (lost sales) and overstocks (costly markdowns) for high-turnover consumer devices, maximizing profitability.

6.6. Optimized Logistics and Transportation:

Agents analyze shipping requirements, carrier availability/rates, traffic, and weather data to optimize routes and modes in SAP TM. They can autonomously book carriers, track shipments, handle exceptions (e.g., delays), and manage documentation.

- Benefit*:* Reduced transportation costs, improved on-time delivery, enhanced visibility.

6.7. Predictive Maintenance & Asset Management:

Agents analyze sensor data (IoT) and maintenance history (SAP Asset Performance Management) to predict equipment failures, automatically create maintenance orders, schedule technicians, and order spare parts.

- Benefit: Reduced downtime, optimized maintenance schedules, extended asset life.

6.8. Enhanced Inventory Management:

Agents monitor stock levels across the network (SAP IBP, EWM), identify imbalances, predict shortages/excess, and trigger automated replenishment orders or inter-site transfers based on cost, lead time, and service level policies [1]. Figure 4, below shows a prototype of the agent in action.

- Benefit*:* Optimized inventory levels, reduced carrying costs, improved fill rates.

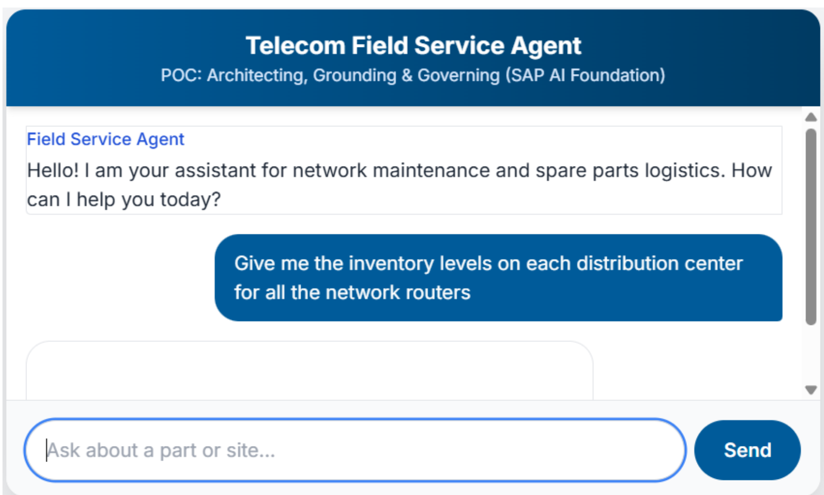

Fig 4: Prototype of Telecom Service AI Agent in action

6.9. Overall Benefits:

Automating complex tasks frees up human resources for strategic activities. Data-driven insights and scenario simulations lead to better, faster decisions. Proactive risk detection and faster response minimize disruption impact [10]. Optimization across procurement, logistics, inventory, and maintenance lowers operational costs. More agile, efficient, and resilient supply chains differentiate businesses. Offering enhanced visibility, reliability, or security as a premium service to customers or partners. Developing specialized agent-driven solutions for niche SCM problems.

7. Approach Methods Implementation Considerations

A successful implementation requires a strategic, phased approach. Align AI agent initiatives with overall business and supply chain objectives. Identify key pain points and high-value opportunities. Secure executive sponsorship. Evaluate data maturity (quality, accessibility, governance), technical infrastructure (BTP readiness), and organizational skills. Identify gaps and plan remediation. Start with pilot projects targeting specific, measurable outcomes (e.g., improving forecast accuracy for a key product line, automating PO creation for specific suppliers). Focus on demonstrating value quickly. Implement SAP AI Ethics guidelines from the start. Define roles, responsibilities, oversight mechanisms, and risk management protocols. Consolidate, cleanse, and structure data required for grounding. Set up data pipelines using SAP Datasphere or SAP Data Intelligence Cloud. Implement vector database capabilities in SAP HANA Cloud.

- Utilize Joule Studio for low-code agent development or leverage SAP AI Core for custom builds.

- Configure GenAI Hub for LLM access, prompt templates, and grounding orchestration.
- Integrate necessary tools and SAP system APIs via BTP Integration Suite.

Build vector indexes in HANA Cloud. Develop RAG pipelines. Integrate Knowledge Graph queries where applicable. Fine-tune retrieval and prompting strategies. Validate agent performance, reliability, security, and ethical alignment in a controlled environment. Test grounding effectiveness and hallucination rates. Simulate various scenarios, including edge cases and potential failures. Deploy the agent for the selected use case with a limited user group. Monitor closely, gather feedback, and iterate. Ensure human oversight mechanisms are effective. Based on pilot success, gradually scale the solution to broader use cases or regions. Continuously monitor performance, retrain models, update grounding data, and optimize configurations. Communicate clearly about the agent's purpose, capabilities, and limitations [11]. Provide training to users on how to interact with and oversee the agents. Redefine roles and processes as needed. The key considerations include focus on iterative value delivery. Ensure agents augment human capabilities and provide intuitive interaction [12]. Embed security throughout the architecture and development lifecycle.

8. Conclusion

Generative AI agents, when thoughtfully architected, reliably grounded, and responsibly governed, represent a significant leap forward for telecom supply chain management. SAP AI Foundation provides a robust and integrated platform for enterprises to harness this potential within their existing SAP landscape. By automating complex decisions, enhancing predictive capabilities, and enabling proactive responses, these agents can transform telecom supply chains from reactive cost centers into resilient, intelligent, and value-driving engines. The journey requires strategic planning, investment in data and skills, and a commitment to ethical principles. However, the benefits – increased efficiency, agility, resilience, and competitiveness – are substantial. By embracing the principles of architecting, grounding, and governing using SAP AI Foundation, organizations can confidently navigate the complexities of modern supply chains and unlock the transformative power of GenAI agents.

References

[1] Jannelli, V., Schoepf, S., Bickel, M., Netland, T., & Brandrup, A. (2024). Agentic LLMs in the Supply Chain: Towards Autonomous Multi-Agent Consensus-Seeking. arXiv preprint arXiv:2411.10184.

[2] Kalaiselvan, S. (2025). Revolutionizing Forecasting with Unified Demand Forecasting for Supply Chain Retail by SAP Customer Activity Repository (CAR) using Machine Learning, Predictive Analysis, and AI. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 6(3), 1-6.

[3] Kraprayoon, J., Williams, Z., & Fayyaz, R. (2025). Ai agent governance: A field guide. arXiv preprint arXiv:2505.21808.

[4] SAP. (2024). SAP AI Ethics Handbook: Applying SAP's Global AI Ethics Policy across the Business AI Lifecycle (Version 2.0). www.sap.com

[5] SAP. (2025). Enabling your own AI Agents with the SAP BTP AI Foundation https://www.sap.com/products/artificial-intelligence/ai-foundation-os.html

[6] Vu, H., Klievtsova, N., Leopold, H., Rinderle-Ma, S., & Kampik, T. (2025, August). Agentic Business Process Management: Practitioner Perspectives on Agent Governance in Business Processes. In International Conference on Business Process Management (pp. 29-43). Cham: Springer Nature Switzerland.

[7] Dignum, V. (2025, May). Responsible AI and Autonomous Agents: Governance, Ethics, and Sustainable Innovation. In Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (pp. 1-2).

[8] Schneider, J., Meske, C., & Kuss, P. (2024). Foundation models: A new paradigm for artificial intelligence. Business & Information Systems Engineering, 66(2), 221-231.

[9] Abbas Khan, M., Khan, H., Omer, M. F., Ullah, I., & Yasir, M. (2024). Impact of artificial intelligence on the global economy and technology advancements. In Artificial General Intelligence (AGI) Security: Smart Applications and Sustainable Technologies (pp. 147-180). Singapore: Springer Nature Singapore.

[10] Borah, A. R., Nischith, T. N., & Gupta, S. (2024, January). Improved learning based on GenAI. In 2024 2nd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT) (pp. 1527-1532). IEEE.

[11] Ferrara, E. (2024). GenAI against humanity: Nefarious applications of generative artificial intelligence and large language models. Journal of Computational Social Science, 7(1), 549-569.

[12] Gu, J., & Yan, Z. (2025). Effects of GenAI interventions on student academic performance: A meta-analysis. Journal of Educational Computing Research, 63(6), 1460-1492.