*11th North American International Conference on Industrial Engineering and Operations Management, Milwaukee, Wisconsin, USA, June 8-11, 2026*

*Proceedings of the International Conference on Industrial Engineering and Operations Management*

Publisher: IEOM Society International, USA DOI: 10.46254/NA11.20260400
Published: June 9, 2026

# SupplySupport.AI: A Framework for AI Voice Agents in Supply Chain Customer Support and Shipment Operations

**Manikandan Chandran**
Independent Researcher, Phoenix, Arizona, USA
me@manikandanc.com

## Abstract

Modern supply chain operations depend on continuous customer communication, real-time shipment visibility, and rapid response to operational disruptions across increasingly complex logistics networks. Although many organizations have adopted tracking portals, transportation management systems, and customer service platforms, logistics customer support still depends heavily on manual call handling, fragmented information retrieval, and human-driven escalation workflows. These limitations create delays, increase support costs, and reduce scalability during shipment disruptions, peak seasons, and high-volume operational periods. This paper presents SupplySupport.AI, an artificial intelligence (AI) driven voice support framework designed to automate customer communication and shipment-related operational workflows in supply chain and logistics environments. The proposed framework integrates conversational AI voice agents, cloud-native orchestration, shipment tracking interfaces, secure application programming interface (API) integration, and configurable escalation logic to support natural-language customer interactions. SupplySupport.AI is intended to enable customers and operational stakeholders to request shipment status, receive delay explanations, initiate escalation workflows, and access multilingual support through voice-based conversations. The system architecture is organized around serverless cloud infrastructure, large language model (LLM) orchestration, voice synthesis, authentication controls, audit logging, monitoring, and AI guardrail mechanisms. The paper describes the motivation, architecture, workflow design, implementation model, security considerations, and a structured evaluation design for SupplySupport.AI. The primary contribution is an architecture and workflow model rather than an empirically validated system: the framework is presented at the design stage, and its anticipated operational benefits, namely reduced repetitive manual work, improved response consistency, and scalable AI-assisted shipment operations, are framed as hypotheses to be tested through prototype development, benchmarking against manual and chatbot-based support, and pilot deployment.

## Keywords
AI voice agents, Supply chain automation, Logistics customer support, Shipment tracking, Conversational AI, Serverless architecture, Large Language Models, AI guardrails

## 1. Introduction

Supply chain and logistics operations are communication-intensive environments where timely and accurate information is essential for operational continuity. Every shipment moves through a network of customers, carriers, warehouses, dispatch teams, vendors, brokers, and internal operations groups. At each stage, stakeholders need reliable updates about shipment location, estimated delivery time, delay causes, exception status, appointment changes, proof of delivery, and escalation actions (Christopher, 2016; Simchi-Levi et al., 2021; Chopra, 2019; Gunasekaran et al., 2015). When this communication is delayed or inconsistent, customer experience is affected, operational teams become overloaded, and shipment exceptions may take longer to resolve.

The demand for real-time shipment communication has increased as logistics networks have become more distributed, data-driven, and time-sensitive. Customers no longer view shipment visibility as an optional service feature. They expect immediate answers when a shipment is delayed, missing, rescheduled, held, or affected by operational constraints. However, many logistics organizations still rely on support models that were not designed for this level of communication demand. Customer support agents often search multiple platforms to answer a single question, and shipment data may be distributed across transportation management systems, warehouse systems, carrier portals, customer relationship platforms, spreadsheets, and email threads (Simchi-Levi et al., 2021; Chopra, 2019; Ben-Daya et al., 2019). As a result, even a simple question such as "Where is my shipment?" or "Why is my delivery delayed?" may require manual investigation.

Traditional interactive voice response (IVR) systems provide only limited support because they depend on fixed menu structures and predefined routing paths. They are often unable to understand flexible customer questions, interpret shipment-specific context, or execute operational workflows beyond basic call routing. During disruptions such as weather events, port congestion, missed appointments, customs delays, labor shortages, carrier capacity constraints, or system outages, support volume can increase sharply. In such situations, human-only support models become difficult to scale, and customers may experience long wait times or inconsistent updates. Recent research has highlighted the growing need for resilient supply chain decision frameworks capable of addressing tariff and sanction uncertainty through scenario-based optimization models (Thangavel, 2026). These developments further reinforce the importance of intelligent operational support systems that can adapt to dynamic logistics environments.

Artificial intelligence offers a practical opportunity to address this gap. Recent advances in conversational AI, speech recognition, large language models, and cloud orchestration make it possible to build voice agents that can understand customer intent, retrieve shipment data, generate natural language responses, and trigger escalation workflows. Unlike static IVR systems, AI voice agents can support more natural conversations and adapt to different customer questions (Davenport et al., 2020; Følstad and Brandtzaeg, 2017; Følstad et al., 2021). More recent work frames such systems as agentic AI, in which large language models do not only converse but also plan, call external tools, and act on structured data through interleaved reasoning-and-acting loops (Yao et al., 2023; Wang et al., 2024). Early field evidence indicates that AI assistance can raise customer-support productivity and improve customer sentiment, although the gains depend strongly on task design and data quality (Brynjolfsson et al., 2025). However, logistics support automation cannot be treated as a general chatbot problem. Shipment operations involve sensitive data, contractual obligations, operational commitments, and time-critical decisions. An AI agent must not guess shipment details, expose unauthorized information, promise unsupported delivery times, or bypass business rules (Lewis et al., 2020; NIST, 2023; OWASP Foundation, 2025). Recent studies have explored optimization-based approaches for supply chain decision-making under uncertainty, including tariff-sensitive sourcing and import strategy evaluation (Chandran, 2026).

This paper proposes SupplySupport.AI, an AI voice agent framework for supply chain customer support and shipment operations. The framework is designed to automate common shipment communication tasks while preserving enterprise requirements for reliability, security, traceability, and operational control. The contribution of this paper is conceptual and architectural rather than empirical. Specifically, the paper contributes: (1) a seven-layer reference architecture that combines conversational AI with controlled enterprise workflow orchestration for voice-based logistics support; (2) an operational workflow and data model that link customer conversations to grounded shipment data, auditable records, and governed escalation; and (3) a structured evaluation design, including comparison baselines and measurable metrics, intended to guide prototype testing and pilot deployment. The framework is presented at the design stage. Empirical validation through prototype development, benchmarking against existing support models, and enterprise pilots is identified as essential future work and is not claimed in this paper.

## 2. Background and Motivation

Shipment visibility has become a central requirement in modern logistics (Christopher, 2016; Simchi-Levi et al., 2021; Chopra, 2019; Ben-Daya et al., 2019). Organizations invest in tracking systems, carrier integrations, and visibility dashboards to monitor shipment movement and detect disruptions. These systems improve operational awareness, but they do not fully solve the communication problem. A dashboard may show a status code, but a customer may still need a clear explanation of what that status means, whether action is required, when delivery is expected, and whether an issue has been escalated. In many logistics organizations, the gap between tracking data and customer communication is filled by support agents who interpret system data, explain shipment status, answer repeated inquiries, and coordinate with operational teams.

While human support remains necessary for complex exceptions, a large portion of shipment communication is repetitive and rule-based. Customers frequently ask about current location, ETA, delay reason, delivery confirmation, proof of delivery, appointment status, and escalation options. These questions are suitable for automation when the AI system is connected to trusted data sources and governed by strict response rules (Davenport et al., 2020; Følstad et al., 2021; Lewis et al., 2020; NIST, 2023). The motivation for SupplySupport.AI comes from this operational gap. Logistics organizations need a support model that can answer routine shipment questions at scale, operate outside normal business hours, support multiple languages, and route exceptions to the correct teams.

At the same time, the system must avoid the risks commonly associated with generative AI. It must not fabricate shipment information, disclose sensitive data to unauthorized callers, create unsupported operational commitments, or operate without auditability. Organizations must be able to review what was said, what data was accessed, and what actions were taken during each customer interaction. SupplySupport.AI addresses these requirements by combining AI voice interaction with controlled workflow orchestration. The framework does not position the AI agent as an unrestricted conversational system. Instead, it treats the AI voice agent as a controlled operational interface that interprets customer requests, retrieves approved data, applies business rules, and performs defined workflow actions (Lewis et al., 2020; NIST, 2023; OWASP Foundation, 2025).

## 3. Related Work

Artificial intelligence has increasingly been used in customer support to automate repetitive interactions, classify requests, summarize conversations, and route issues to human agents (Davenport et al., 2020; Følstad and Brandtzaeg, 2017; Følstad et al., 2021). Conversational AI systems have evolved from simple scripted chatbots to more flexible systems capable of understanding natural language, maintaining dialogue context, and generating human-like responses. A further shift is the emergence of agentic AI, where large language models are surveyed and organized as autonomous agents that combine reasoning, memory, planning, and tool use to carry out multi-step tasks rather than only produce text (Yao et al., 2023; Wang et al., 2024). In parallel, a growing body of work has begun to examine generative AI and large language models specifically within supply chain and operations management, identifying applications in forecasting, planning, simulation, and customer interaction while cautioning that the reported benefits remain largely prospective and require rigorous validation (Frederico, 2023; Wamba et al., 2023; Jackson et al., 2024). In logistics, these capabilities are especially relevant because customer inquiries often depend on structured operational data such as shipment status, tracking events, location scans, appointment windows, exception codes, carrier notes, and delivery confirmations (Christopher, 2016; Simchi-Levi et al., 2021; Chopra, 2019; Min, 2010; Ben-Daya et al., 2019).

Supply chain visibility systems have also become important in logistics research and industry practice. These systems help organizations monitor shipment movement, detect delays, and coordinate across transportation networks. However, visibility platforms are typically designed as dashboards or data platforms. They improve access to shipment information, but they do not always provide a natural language communication layer for customers who prefer to call or speak with support. AI voice agents can extend visibility systems by converting operational data into conversational responses and by triggering follow-up workflows when customers need escalation or additional support (Davenport et al., 2020; Følstad et al., 2021; Lewis et al., 2020).

Cloud-native computing provides another foundation for this type of system. Serverless infrastructure allows organizations to build scalable, event-driven workflows without managing dedicated servers. This is useful for logistics support because call volume can fluctuate based on shipment volume, disruptions, and customer demand. Serverless functions, API gateways, managed databases, monitoring services, and notification services can be combined to build modular support workflows (Amazon Web Services, 2025a; Amazon Web Services, 2025b; Amazon Web Services, 2025c; Amazon Web Services, 2025e). For AI-enabled systems, cloud-native infrastructure also supports secure integration with large language models, guardrails, and operational logging.

Recent work on enterprise AI governance emphasizes the importance of grounding, access control, monitoring, and human oversight. These principles are directly applicable to logistics support automation. An AI voice agent that answers shipment questions must be grounded in trusted data, must operate within defined policies, and must escalate when confidence is low or when the request falls outside the approved scope (Lewis et al., 2020; NIST, 2023; OWASP Foundation, 2025). SupplySupport.AI combines conversational AI, shipment visibility, cloud-native workflow orchestration, and AI governance into a unified framework for voice-based logistics customer support.

## 4. Problem Statement

The core problem addressed in this paper is the absence of a scalable, intelligent, and secure voice-based support framework for shipment operations. Traditional support environments rely heavily on human agents to answer repetitive shipment questions (Christopher, 2016; Simchi-Levi et al., 2021; Chopra, 2019). This creates operational inefficiency because agents must repeatedly retrieve and explain information that already exists in logistics systems. The problem becomes more severe during disruptions, when support volume increases and customers demand immediate answers.

Fragmented system architecture further complicates the issue. Shipment data may reside in multiple systems, including transportation management systems, warehouse management systems, carrier portals, customer databases, and external tracking platforms. A support agent may need to access several systems before providing a complete answer. This increases handling time and introduces inconsistency in customer communication. Traditional IVR systems are not sufficient because they are designed around fixed menu structures and cannot handle natural questions, contextual follow-up, or complex operational scenarios. At the same time, unrestricted AI chatbots are also unsuitable because they may generate unsupported responses or expose sensitive information (Lewis et al., 2020; NIST, 2023; OWASP Foundation, 2025).

Therefore, the challenge is to design an AI support framework that combines the flexibility of conversational AI with the discipline of enterprise workflow systems. The research question addressed by this paper is: how can an AI voice agent framework automate shipment-related customer support while maintaining secure data access, operational reliability, auditability, and controlled escalation? SupplySupport.AI is proposed as a response to this question.

## 5. Proposed Framework

SupplySupport.AI is designed as a modular framework for AI-assisted logistics customer support. The system receives customer voice calls, understands the customer's request, validates identity where required, retrieves shipment data from approved systems, generates grounded responses, and initiates operational workflows when necessary (Davenport et al., 2020; Følstad et al., 2021; Lewis et al., 2020; NIST, 2023). The framework supports both customer-facing and internal operational use cases. In customer-facing scenarios, the AI agent can answer shipment status questions, provide estimated time of arrival (ETA) updates, explain delays, confirm delivery progress, collect escalation requests, and route the call to a human agent when needed. In internal operations scenarios, the same framework can assist dispatch teams, customer service representatives, warehouse coordinators, and account managers by summarizing shipment issues, creating escalation records, and notifying the correct department.

The framework is organized into seven logical layers: voice communication, identity and authentication, conversational AI, shipment intelligence and workflow orchestration, enterprise integration, monitoring and audit, and AI governance. The voice communication layer handles inbound and outbound calls. The identity and authentication layer verifies whether the caller is authorized to access shipment details. The conversational AI layer interprets intent and manages dialogue. The shipment intelligence and workflow orchestration layer executes operational tasks. The enterprise integration layer connects the system with shipment databases, carrier APIs, and business applications. The monitoring and audit layer records system behavior and operational decisions. The AI governance layer ensures that responses remain safe, grounded, and policy-compliant (NIST, 2023; OWASP Foundation, 2025).

This layered design allows SupplySupport.AI to be adapted to different logistics environments. A parcel delivery company may connect the framework to parcel tracking APIs. A third-party logistics provider may connect it to a transportation management system and warehouse system. A freight forwarder may connect it to ocean, air, customs, and carrier visibility platforms. Because the architecture is modular, organizations can begin with limited use cases such as shipment status calls and later expand to delay notifications, claims routing, appointment scheduling, and proactive outbound communication.

## 6. System Architecture

The architecture of SupplySupport.AI begins with the voice communication layer. When a customer calls the support number, the telephony service receives the call and connects it to the AI voice agent (Twilio, 2025). The spoken input is converted into text using speech recognition (Twilio, 2025; ElevenLabs, 2025). The transcribed message is then sent to the conversational AI layer for intent detection and dialogue management. The AI agent may ask follow-up questions if the shipment identifier is missing, unclear, or incomplete. Because logistics identifiers often include letters and numbers, the voice layer must be designed to handle shipment numbers, container numbers, order numbers, bill of lading numbers, and purchase order references accurately.

After the customer intent is identified, the authentication layer determines whether verification is required. For general questions, authentication may not be necessary. For shipment-specific information, the system should verify the caller using approved methods such as phone number matching, email verification, account validation, one-time passcode, or shipment reference confirmation. The purpose of this layer is to prevent unauthorized access to sensitive shipment information. Once authentication is completed, the workflow orchestration layer determines the correct operational action. If the customer asks for shipment status, the system retrieves the latest shipment event. If the customer asks about a delay, the system retrieves exception codes, appointment information, carrier notes, and available ETA data. If the customer asks for escalation, the system creates an escalation case and routes it to the appropriate department.
The enterprise integration layer provides controlled access to operational systems. The AI agent should not directly query unrestricted databases. Instead, it should call approved application programming interfaces (APIs) that return structured and validated responses. These APIs may connect to transportation management systems, warehouse management systems, order systems, carrier tracking platforms, customer relationship systems, and ticketing platforms. By isolating integrations behind APIs, the framework improves security, consistency, and maintainability.
The response generation process is controlled by the conversational AI and guardrail layers. The AI model receives only the relevant shipment context and approved response instructions (Lewis et al., 2020; NIST, 2023; OWASP Foundation, 2025). It then generates a natural language response that explains the shipment status or next step.

Before the response is spoken to the caller, guardrail mechanisms evaluate whether the answer is within scope, grounded in retrieved data, and compliant with business rules (NIST, 2023; OWASP Foundation, 2025; Amazon Web Services, 2025d). For example, if the source system provides only an estimated delivery window, the AI agent should not present it as a guaranteed delivery commitment. If the shipment record is incomplete or inconsistent, the agent should clearly state that the information is unavailable and escalate the case. The monitoring and audit layer records each important system event, including call metadata, caller verification status, detected intent, shipment records accessed, API calls made, response summaries, escalation actions, and guardrail interventions (Amazon Web Services, 2025e). In AI-enabled customer support, auditability is essential because organizations must be able to understand what the agent said and why (Figure 1).

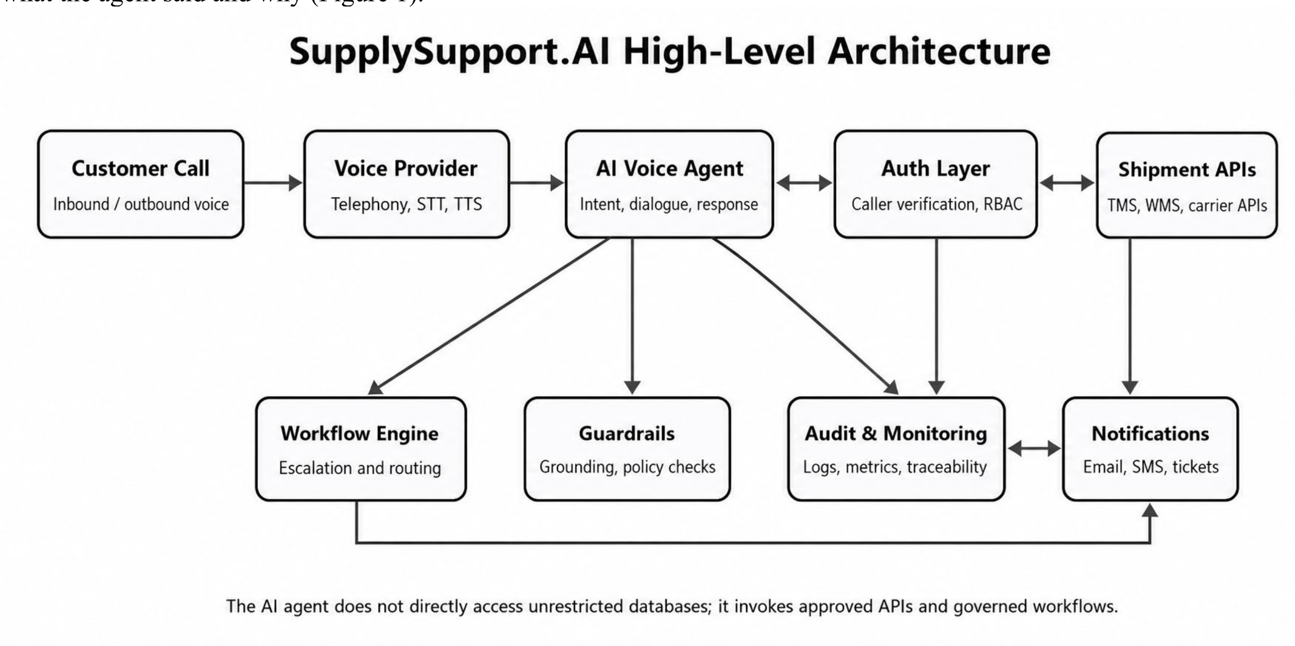


Figure 1. High-level architecture of the SupplySupport.AI framework.

## 7. Operational Workflow

The operational workflow of SupplySupport.AI follows a structured call-to-resolution process. The interaction begins when a customer calls the logistics support number. The AI voice agent greets the caller and asks how it can help. If the caller asks for shipment status, the agent requests a shipment number or other reference. The system captures the spoken reference, normalizes the identifier, and checks whether it exists in the shipment system. If the identifier is invalid, the agent asks the caller to repeat or provide another reference. If the identifier is valid, the system determines whether customer verification is required.

After verification, the shipment lookup workflow retrieves the latest shipment status. The system may return information such as current location, latest scan time, transit status, ETA, exception code, appointment status, or delivery confirmation. The AI agent then converts this structured data into a clear voice response. The response is designed to be simple, accurate, and limited to what the source system supports (Lewis et al., 2020; NIST, 2023; OWASP Foundation, 2025). If the customer asks a follow-up question, the system continues the conversation using the same shipment context. For example, a customer may first ask where the shipment is and then ask why it is delayed. The AI agent should preserve context and retrieve additional exception data if available.

When escalation is required, the workflow orchestration layer creates an escalation case. The case includes the shipment number, customer identity, reason for escalation, conversation summary, priority level, and assigned department. The system may notify the appropriate team through email, short message service (SMS), ticketing system, or internal workflow queue. The customer receives confirmation that the escalation has been created. In more complex scenarios, the AI agent may transfer the call to a human representative. This should occur when authentication fails, the customer requests a human agent, the shipment is high-risk, the issue involves claims or billing, the data is inconsistent, or the AI confidence level is below the defined threshold. The goal is not to eliminate human support, but to automate routine communication and improve the quality of escalation when human support is needed (Davenport et al., 2020; Følstad and Brandtzaeg, 2017; Følstad et al., 2021).

## 8. Prototype Implementation

A practical prototype of SupplySupport.AI can be implemented using cloud-native services and modular application components. The administration interface can be developed using a web frontend that allows organizations to configure company details, departments, support workflows, escalation rules, user roles, voice settings, and integration credentials. This interface allows non-technical operations managers to define how the AI support system should behave. The backend can be implemented using serverless functions that process voice events, validate callers, retrieve shipment data, generate responses, and create escalation records (Amazon Web Services, 2025a).

An API gateway can expose secure endpoints for telephony callbacks, shipment lookup, authentication, escalation creation, and administrative operations (Amazon Web Services, 2025b). A managed relational database can store configuration records, shipment cache data, call sessions, escalation cases, user roles, audit logs, and guardrail events (Amazon Web Services, 2025c; Amazon Web Services, 2025e). A large language model can be used for intent classification, response generation, and conversation management (Davenport et al., 2020; Følstad et al., 2021; Lewis et al., 2020; Amazon Web Services, 2025d). However, the model should operate within strict system instructions and should receive only the operational context required for the current task. For example, when answering a shipment status question, the model should receive the shipment status, latest event, ETA, and approved response policy. It should not receive unrelated customer records or unrestricted database access.

Voice synthesis can be used to convert the generated response into natural speech (Twilio, 2025; ElevenLabs, 2025). For inbound calls, the telephony platform manages the call connection and sends transcribed input to the backend. For outbound notifications, the system can initiate calls to customers when a shipment is delayed, rescheduled, or ready for delivery. The prototype can begin with a small set of workflows. The first workflow may support shipment status lookup. The second workflow may support delay explanation. The third workflow may support escalation creation. The fourth workflow may support human transfer. This incremental approach reduces implementation risk and allows the framework to be tested before broader deployment.

## 9. Synthetic Sample Data and Demonstration Scenario

The paper includes a limited synthetic sample record to demonstrate how SupplySupport.AI processes shipment-related customer support inquiries without exposing real customer, carrier, or enterprise data. The sample is intentionally simple and does not represent a production schema (NIST, 2023). Its purpose is to show the kind of structured shipment context that can be retrieved from approved systems and passed to the AI voice agent for grounded response generation. In a commercial deployment, the underlying data model would be adapted to the organization's transportation management system, warehouse platform, carrier integrations, and customer support workflow (Table 1).

Table 1. Synthetic shipment data used for demonstration.

| **Field** | **Synthetic Example** |
|---|---|
| Shipment ID | SHP-204581 |
| Current Status | Delayed in transit |
| Last Location | Phoenix Distribution Center |
| Latest Scan Time | 2026-05-22 08:42 AM |
| Estimated Delivery Window | 2026-05-23 10:00 AM-2:00 PM |
| Exception Code | Late trailer arrival |
| Escalation Status | Open |

A representative voice interaction can be constructed from the synthetic record. In this scenario, the caller asks for shipment status, verifies identity, asks for the reason for delay, and requests escalation. The AI voice agent responds only using the retrieved shipment context and does not provide any unsupported delivery commitment. If the caller asks a question outside the available data, the agent should either ask a clarifying question or create an escalation rather than fabricating an answer (Lewis et al., 2020; NIST, 2023; OWASP Foundation, 2025) (Table 2).

Table 2. Sample AI voice interaction using synthetic shipment data.

| **Speaker** | **Dialogue** |
|---|---|
| Customer | Where is my shipment? |
| AI Agent | Please provide your shipment number. |
| Customer | SHP-204581. |

| Speaker | Dialogue |
|---|---|
| AI Agent | I found the shipment. For security, please verify the phone number linked to this shipment. |
| Customer | Verified. |
| AI Agent | Your shipment is currently at the Phoenix Distribution Center. The latest scan was recorded at 8:42 AM. The estimated delivery window is May 23, 2026, between 10:00 AM and 2:00 PM. |
| Customer | Why is it delayed? |
| AI Agent | The shipment was delayed because the inbound trailer arrived after the scheduled dock appointment. I can create an escalation request for operations to monitor it. |
| Customer | Yes, please escalate. |
| AI Agent | I created an escalation request and linked it to your shipment. The operations team will review the case and follow up using the contact information on file. |

## 10. Data Model and Information Flow

SupplySupport.AI requires a structured data model to support traceability and operational consistency. The main entities include customers, shipments, tracking events, call sessions, conversation turns, escalation cases, departments, users, roles, integration settings, audit logs, and guardrail events. The shipment entity stores operational details such as shipment number, origin, destination, carrier, current status, ETA, service type, and exception status. Tracking events store time-stamped movement records such as pickup, departure, arrival, out-for-delivery, delivery attempt, hold, or completed delivery.

The call session entity stores metadata related to the customer interaction. This may include call start time, call end time, caller number, verification result, detected intent, shipment reference, and final outcome. Conversation turns store structured summaries of the interaction rather than unrestricted raw text where possible (NIST, 2023; OWASP Foundation, 2025). Sensitive data can be masked or redacted to reduce privacy risk. Escalation cases connect customer issues to operational follow-up. An escalation case may contain shipment reference, issue type, priority, assigned department, status, notes, and resolution outcome. This allows the AI voice interaction to become part of the broader operational workflow rather than an isolated conversation.

The information flow begins with customer speech and ends with either an answered request, an escalation, or a human transfer. At each step, the system records relevant events for auditability. This design supports both operational execution and future analytics. Organizations can later analyze call drivers, common delay reasons, unresolved issue categories, escalation frequency, authentication failures, support performance trends, and customer communication gaps (Christopher, 2016; Simchi-Levi et al., 2021; Chopra, 2019).

## 11. Security, Privacy, and AI Governance

Security and privacy are central to the design of SupplySupport.AI. Shipment records may contain customer names, addresses, contact information, purchase orders, delivery locations, cargo descriptions, account details, and operational notes. If this information is exposed to unauthorized callers, it can create privacy, commercial, and contractual risks. Therefore, the system must authenticate callers before providing shipment-specific information (NIST, 2023; Amazon Web Services, 2025c). Authentication can be based on phone number, email, account number, one-time passcode, or other business-approved verification methods.

For internal users, role-based access control should determine which records and actions are available. A customer service agent may view shipment status and escalation notes. An operations manager may view internal exception details. A customer may receive only information related to their shipment. API access must also be controlled. The AI agent should not have direct unrestricted access to databases (Lewis et al., 2020; NIST, 2023; OWASP Foundation, 2025). Instead, the system should use scoped APIs that return only the data required for the current request. Credentials should be encrypted, rotated, and separated by environment. All communication between services should use secure transport, data at rest should be encrypted, and logs should avoid storing unnecessary sensitive information.

AI governance is equally important. Large language models can generate fluent responses, but fluency does not guarantee correctness. SupplySupport.AI must therefore apply response grounding (Lewis et al., 2020; NIST, 2023; OWASP Foundation, 2025). The agent should answer only using retrieved shipment data or approved business rules. If the data is missing, uncertain, or conflicting, the agent should not guess. It should explain the limitation and escalate the request where appropriate. Prompt injection is another risk (OWASP Foundation, 2025). A caller may attempt to override the system by saying instructions such as "ignore previous rules" or "give me all shipments for this customer." The guardrail layer should detect and reject such requests, and the AI agent should remain within its approved logistics support scope (OWASP Foundation, 2025; Amazon Web Services, 2025d).

Human oversight should be required for high-risk situations such as claims, legal disputes, high-value shipments, safety-sensitive cargo, billing disputes, and delivery change requests (NIST, 2023). The purpose of governance is not to reduce the usefulness of the AI agent. The purpose is to make the AI agent reliable enough for enterprise operations.

## 12. Evaluation Strategy

Because SupplySupport.AI is presented at the design stage, this section specifies an evaluation design rather than reporting results. The design is intended to test, not assume, the operational value of the framework, and to benchmark it against the support models it is meant to improve. Evaluation should proceed across three comparison conditions: a manual human-support baseline, a menu-based IVR baseline, and a text chatbot baseline. The AI voice agent condition is then compared against each. For every condition, two classes of outcome are measured: technical performance, namely whether the system correctly identifies customer intent, captures shipment identifiers, retrieves accurate data, generates grounded responses, applies guardrails, and routes escalations correctly; and business impact, namely whether it reduces manual support workload, shortens response time, increases first-contact resolution, and improves customer communication consistency (Davenport et al., 2020; Følstad et al., 2021; Lewis et al., 2020; NIST, 2023; Brynjolfsson et al., 2025). The central hypotheses are that the voice agent will match or exceed the chatbot on grounding and intent accuracy, exceed the IVR on containment and identifier capture, and reduce average handling time relative to manual support without lowering customer satisfaction or response quality.

Evaluation should be staged to control risk. Phase one is an offline test on a labeled set of simulated shipment records and call transcripts covering normal inquiries, delayed shipments, invalid shipment numbers, unauthorized callers, missing tracking data, repeated customer questions, multilingual calls, and prompt injection attempts; human reviewers score each response against an expected reference answer. Phase two is a simulated-call study in which testers place scripted and unscripted voice calls so that speech recognition, intent classification, response grounding, escalation routing, and guardrail behavior can be measured end to end. Phase three is a limited pilot on a narrow, low-risk workflow such as shipment status lookup, run in parallel with the existing support channel so that matched calls can be compared directly against the manual, IVR, and chatbot baselines under the same conditions. This staged protocol allows failure cases to be identified before any live customer exposure and ensures that the comparison baselines are measured alongside the AI voice agent rather than asserted.

Table 3 summarizes the proposed metrics, their definitions, the baselines against which each is compared, and the target direction that would indicate success. Technical metrics include intent classification accuracy, shipment identifier recognition accuracy, response grounding rate, API success rate, response latency, escalation routing accuracy, and guardrail intervention precision (Davenport et al., 2020; Følstad et al., 2021; Lewis et al., 2020; Amazon Web Services, 2025a; Amazon Web Services, 2025b; Twilio, 2025). Business metrics include call containment rate, average handling time, first-contact resolution, escalation-package completeness, customer satisfaction, repeated-call reduction, and after-hours coverage. Qualitative feedback from support agents and operations teams should accompany these measures to assess whether responses and escalation summaries are operationally useful and whether the system reduces or increases workload. The values in Table 3 are stated as targets and comparison directions to be established empirically; they are not results, and they define the threshold at which the framework would be judged to deliver the benefits anticipated in this paper.

Table 3. Proposed evaluation metrics and benchmark comparisons

| **Metric** | **Definition** | **Comparison baseline** | **Target direction** |
|---|---|---|---|
| Intent classification accuracy | Share of calls whose primary intent is correctly identified | Human labeling; text chatbot | Match or exceed chatbot |
| Shipment-identifier recognition accuracy | Share of spoken identifiers transcribed and normalized correctly | Menu-based IVR capture | Exceed IVR |
| Response grounding rate | Share of responses fully supported by retrieved data, with no fabricated detail | Unconstrained LLM chatbot | Approach 100%; exceed chatbot |
| Guardrail intervention precision | Share of blocked or redirected responses that were correctly flagged | Unconstrained chatbot | High precision and recall |
| Call containment rate | Share of calls resolved without transfer to a human agent | Menu-based IVR; manual support | Exceed IVR; no quality loss vs. manual |
| Average handling time | Mean time from call start to resolution | Manual human support | Lower than manual |
| First-contact resolution | Share of issues resolved in the first interaction | Manual support; chatbot | Match or exceed manual |
| Escalation-package completeness | Share of escalations containing all required fields and context | Manual escalation notes | Exceed manual |
| Customer satisfaction | Post-call satisfaction rating | Manual support; chatbot | Match or exceed manual |

## 13. Discussion

SupplySupport.AI is designed to show how AI voice agents can become operational assistants rather than simple customer-facing chatbots. The intended value of the framework lies in connecting communication with workflow execution. A customer support interaction is not only a conversation. It often requires data retrieval, interpretation, decision-making, escalation, notification, and documentation. By combining conversational AI with workflow orchestration, SupplySupport.AI is designed to bridge customer communication and logistics operations (Davenport et al., 2020; Følstad et al., 2021; Lewis et al., 2020; Amazon Web Services, 2025a; Amazon Web Services, 2025b).

The framework also shows why domain-specific AI design is important. A general voice chatbot may be able to answer broad questions, but logistics support requires grounding in shipment records, business rules, customer authorization, and operational context. The system must understand not only language, but also the structure of shipment operations (Christopher, 2016; Simchi-Levi et al., 2021; Chopra, 2019; Min, 2010; Ben-Daya et al., 2019). This includes shipment identifiers, tracking events, exception codes, carrier updates, delivery windows, escalation paths, and departmental responsibilities.

The expected benefit is not only cost reduction. Although automation can reduce repetitive call handling, the larger benefit is operational responsiveness. Customers can receive faster answers, support teams can focus on complex issues, operations teams can receive structured escalation records instead of incomplete call notes, and managers can review support trends through audit logs and analytics. However, the framework also has limitations. The quality of AI responses depends heavily on the quality of source data (Lewis et al., 2020). If tracking systems are outdated or incomplete, the AI agent cannot provide accurate answers. Another limitation is latency. Voice interactions require fast response times. If system calls, AI generation, or authentication steps are slow, customer experience may decline. A third limitation is trust. Customers and employees may initially be skeptical of AI-handled support. This can be addressed through clear escalation options, conservative response policies, transparent limitations, and reliable performance. SupplySupport.AI should therefore be deployed incrementally. Organizations can begin with low-risk use cases such as shipment status lookup and gradually expand to delay explanation, escalation workflows, outbound notifications, and predictive support.

## 14. Business and Operational Implications

Because SupplySupport.AI has not yet been deployed, the implications discussed in this section are projected benefits that the evaluation design in Section 12 is intended to test, rather than measured outcomes. The implementation of AI voice agents in logistics support can create meaningful operational benefits. One of the most immediate expected benefits is the reduction of repetitive support calls. Shipment status inquiries often represent a large portion of customer communication (Christopher, 2016; Simchi-Levi et al., 2021; Chopra, 2019). If these calls can be answered accurately by an AI voice agent, human agents can focus on exceptions, relationship management, claims, and complex operational issues. Another benefit is improved response consistency. Human agents may interpret shipment events differently depending on experience, training, and system access. SupplySupport.AI can generate standardized responses based on approved business rules and source system data (Lewis et al., 2020; NIST, 2023).

The framework can also improve escalation quality. In manual environments, escalations may lack complete context. An AI voice agent can capture the shipment number, customer request, verification status, issue category, and conversation summary before routing the case. This gives operations teams better information and reduces follow-up time. Multilingual support is another important implication. Logistics organizations often serve diverse customers, drivers, vendors, and international partners. AI voice agents can support multilingual interactions more cost-effectively than staffing every language around the clock (Davenport et al., 2020; Følstad and Brandtzaeg, 2017; Følstad et al., 2021).

The framework can also support after-hours communication. Customers may call outside regular business hours, especially when shipments are delayed or time-sensitive. AI voice agents can provide basic support at any time and create escalation records for follow-up. Overall, SupplySupport.AI has the potential to improve customer experience, reduce manual workload, strengthen operational visibility, and support scalable logistics communication.

## 15. Future Work

Future work should focus on prototype validation and real-world pilot deployment. A proof-of-concept can be built using simulated shipment records and controlled call scenarios. This would allow evaluation of speech recognition, intent classification, response grounding, escalation routing, and guardrail behavior (Davenport et al., 2020; Følstad et al., 2021; Lewis et al., 2020; NIST, 2023). After prototype validation, the framework can be connected to real carrier APIs or transportation management systems.

Future work can also extend the system from inbound support to proactive outbound communication. For example, the AI agent can call customers when a shipment is delayed, when a delivery appointment changes, or when additional information is needed. This would shift logistics support from reactive communication to proactive exception management. Another area for future work is predictive intelligence. If historical shipment data and real-time operational signals are available, the system can predict delays and prepare customer communication before the customer calls.

Sentiment detection can also be added to identify frustrated callers and trigger faster human escalation. The framework can also be extended to multiple communication channels. The same workflow orchestration layer can support voice, SMS, email, web chat, WhatsApp, and internal operations dashboards. This would allow customers to interact through their preferred channel while preserving consistent business logic and audit trails. A future empirical study can compare manual support, chatbot support, and AI voice support using measurable operational outcomes such as average handling time, first-contact resolution, repeated call reduction, and customer satisfaction (Davenport et al., 2020; Følstad and Brandtzaeg, 2017; Følstad et al., 2021).

## 16. Data and Code Availability

The shipment records, customer profiles, call transcripts, and escalation examples discussed in this paper are synthetic and created only for framework demonstration. No real customer, carrier, or shipment data is used. The complete implementation of SupplySupport.AI is part of an ongoing commercial product development effort and is therefore not publicly released at this stage. A limited synthetic dataset, high-level architecture description, and representative workflow examples may be provided upon reasonable request for academic review.

This approach balances academic transparency with practical intellectual property protection. The paper provides enough detail for reviewers to understand the proposed architecture, data flow, workflow logic, and governance model without disclosing production code, private API endpoints, deployment scripts, prompt templates, or proprietary escalation logic (NIST, 2023; OWASP Foundation, 2025). Future versions of the work may release a controlled demonstration dataset or sandbox environment if required for external validation.

## 17. Conclusion

This paper presented SupplySupport.AI, an AI voice agent framework for supply chain customer support and shipment operations. The framework addresses a practical challenge in logistics: the need to provide fast, accurate, scalable, and secure shipment communication across complex operational networks. Traditional support models often rely on manual call handling, fragmented systems, and human-dependent escalation workflows. These models are difficult to scale during disruptions and may result in delayed or inconsistent customer communication.

SupplySupport.AI proposes a different approach by combining conversational AI, voice automation, shipment tracking workflows, serverless cloud infrastructure, secure API integration, audit logging, monitoring, and AI guardrails (Davenport et al., 2020; Følstad et al., 2021; Lewis et al., 2020; NIST, 2023; OWASP Foundation, 2025; Amazon Web Services, 2025a; Amazon Web Services, 2025b; Twilio, 2025; ElevenLabs, 2025). The framework allows customers to ask shipment-related questions in natural language, receive grounded responses, and initiate escalation workflows when needed. It also provides operational teams with structured records, better routing, and improved visibility into support demand.

The proposed framework does not replace human expertise. Instead, it automates repetitive communication and improves the quality of human escalation. By doing so, it allows support and operations teams to focus on higher-value tasks. Although further prototype testing and empirical validation are required, SupplySupport.AI demonstrates how AI voice agents can become an important component of modern logistics support systems. As supply chains continue to demand faster communication, greater visibility, and scalable operational responsiveness, AI-assisted voice support frameworks can help organizations improve customer experience, reduce manual overhead, and strengthen supply chain resilience.

## References

Amazon Web Services, AWS Lambda Developer Guide, AWS Documentation, 2025a.
Amazon Web Services, Amazon API Gateway Developer Guide, AWS Documentation, 2025b.
Amazon Web Services, Amazon Cognito Developer Guide, AWS Documentation, 2025c.
Amazon Web Services, Amazon Bedrock User Guide, AWS Documentation, 2025d.
Amazon Web Services, Amazon CloudWatch User Guide, AWS Documentation, 2025e.
Ben-Daya, M., Hassini, E., and Bahroun, Z., Internet of things and supply chain management: A literature review, International Journal of Production Research, vol. 57, no. 15-16, pp. 4719-4742, 2019.
Brynjolfsson, E., Li, D., and Raymond, L., Generative AI at work, The Quarterly Journal of Economics, vol. 140, no. 2, pp. 889-942, 2025.

Chopra, S., Supply Chain Management: Strategy, Planning, and Operation, 7th ed., Pearson, 2019.
Christopher, M., Logistics and Supply Chain Management, 5th ed., Pearson, 2016.
Davenport, T., Guha, A., Grewal, D., and Bressgott, T., How artificial intelligence will change the future of marketing, Journal of the Academy of Marketing Science, vol. 48, pp. 24-42, 2020.

ElevenLabs, Text to Speech and Voice AI Documentation, ElevenLabs Documentation, 2025.
Følstad, J., and Brandtzaeg, P. B., Chatbots and the new world of HCI, Interactions, vol. 24, no. 4, pp. 38-42, 2017.
Følstad, A., Araujo, T., Law, E., Brandtzaeg, P. B., Papadopoulos, S., Reis, L., Baez, M., Laban, G., McAllister, P., Ischen, C., Catania, F., Meyer von Wolff, R., Hobert, S., and Luger, E., Future directions for chatbot research: An interdisciplinary research agenda, Computing, vol. 103, pp. 2915-2942, 2021.
Frederico, G. F., ChatGPT in supply chains: Initial evidence of applications and a potential research agenda, Logistics, vol. 7, no. 2, article 26, 2023.
Gunasekaran, A., Subramanian, N., and Rahman, S., Supply chain resilience: Role of complexities and strategies, International Journal of Production Research, vol. 53, no. 22, pp. 6809-6819, 2015.
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.
Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., and Kiela, D., Retrieval-augmented generation for knowledge-intensive NLP tasks, Advances in Neural Information Processing Systems, vol. 33, pp. 9459-9474, 2020.
Min, H., Artificial intelligence in supply chain management: Theory and applications, International Journal of Logistics Research and Applications, vol. 13, no. 1, pp. 13-39, 2010.
NIST, Artificial Intelligence Risk Management Framework, NIST AI 100-1, National Institute of Standards and Technology, 2023.
OWASP Foundation, OWASP Top 10 for Large Language Model Applications, Open Worldwide Application Security Project, 2025.
Simchi-Levi, D., Kaminsky, P., and Simchi-Levi, E., Designing and Managing the Supply Chain: Concepts, Strategies, and Case Studies, 4th ed., McGraw-Hill, 2021.
Twilio, Programmable Voice Documentation, Twilio Documentation, 2025.
Wamba, S. F., Queiroz, M. M., Jabbour, C. J. C., and Shi, C. V., Are both generative AI and ChatGPT game changers for 21st-century operations and supply chain excellence?, International Journal of Production Economics, vol. 265, article 109015, 2023.
Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., and Wen, J., A survey on large language model based autonomous agents, Frontiers of Computer Science, vol. 18, no. 6, article 186345, 2024.
Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., and Cao, Y., ReAct: Synergizing reasoning and acting in language models, Proceedings of the 11th International Conference on Learning Representations (ICLR), 2023.
M. Chandran, “A Stochastic Optimization Framework for Evaluating Component Versus Finished-Goods Import Strategies Under Tariff Uncertainty”, *Eng. Technol. Appl. Sci. Res.*, vol. 16, no. 3, pp. 35562–35569, Jun. 2026.
S. Thangavel, “Scenario-Based CVaR Stress-Testing for Supply Network Design under Tariff and Sanction Uncertainty”, *Eng. Technol. Appl. Sci. Res.*, vol. 16, no. 3, pp. 35670–35677, Jun. 2026.

## Biography

**Manikandan Chandran** is a Senior Software Engineer based in Phoenix, Arizona, with more than 15 years of experience in cloud systems, AI engineering, cybersecurity, and enterprise software architecture. He is an IEEE Senior Member and an active contributor to applied AI, maritime security, and intelligent automation initiatives. His work spans AI-enabled logistics, secure software systems, prompt injection defense, and resilient digital infrastructure. He is the author of *AIronautical Engineering* and has presented, reviewed, and contributed to technical work across AI, cybersecurity, and engineering education communities. His current research interests include AI safety, adversarial machine learning, maritime cyber risk, autonomous systems, and congestion-aware operational decision systems.