Research Article *Open Access*

Journal of Pioneering Artificial Intelligence Research

ISSN: 3069-0846

DOI: doi.org/10.63721/26JPAIR0156

NAUTILUS: An Agentic Architecture for Real-Time Supply Chain Intelligence

Anubhav Tewari

Warren, Michigan, USA

Citation: *Anubhav Tewari (2026) NAUTILUS: An Agentic Architecture for Real-Time Supply Chain Intelligence. J of Poin Artf Research 2(5), 1-7. WMJ-JPAIR-156*

Abstract

Global supply chains are increasingly exposed to compounding disruptions spanning geopolitical events, commodity price volatility, natural disasters, trade policy shifts, and regulatory compliance risk. Existing enterprise procurement tools address these domains in isolation, requiring analysts to manually synthesis fragmented signals across disconnected systems. This paper presents NAUTILUS Terminal, an agentic supply chain intelligence architecture that unifies real-time multi-source data ingestion, large language model (LLM)-assisted classification and recommendation, and structured risk fusion into a single decision-support interface. The system decomposes procurement intelligence into seven autonomous API agents covering commodity pricing, geopolitical incidents, seismic events, wildfire propagation, maritime port status, foreign exchange dynamics, and trade news, whose outputs are synthesized by a central reasoning agent into actionable sourcing directives, risk scores, and compliance assessments. We describe the system architecture, data integration methodology, LLM orchestration strategy, and evaluate the system across 40 commodity categories. NAUTILUS demonstrates that agentic composition of heterogeneous live data streams, mediated by structured prompting and deterministic reasoning, produces procurement intelligence that is broader in coverage and lower in latency than conventional enterprise solutions.

***Corresponding author**:Anubhav Tewari, Warren, Michigan, USA.

Received: 16.09.2026 **Accepted**: 21.09.2026 **Published**: 26.10.2026

Keywords: Supply Chain Intelligence, Agentic AI, Procurement Decision Support, Multi-Source Risk Fusion, Large Language Models, Real-Time Systems

Introduction

The fragility of global supply chains has become a defining operational challenge of the 2020s. The COVID-19 pandemic, the Red Sea shipping crisis, escalating trade tensions between major economic blocs, and the accelerating frequency of extreme weather events have collectively exposed the inadequacy of reactive, siloed procurement intelligence. Organisations that rely on quarterly analyst re-ports, static supplier databases, and manually curated risk dashboards are structurally disadvantaged relative to the pace at which supply disruptions now evolve.

Contemporary enterprise procurement platforms, including SAP Ariba, Coupa, and Jaggaer, provide workflow automation and spend analytics but offer limited capability for real-time geopolitical or commodity risk assessment. Specialised risk platforms such as Resilinc and Everstream Analytics provide supplier-level risk monitoring but require extensive onboarding, carry enterprise licensing costs, and do not expose the underlying data signals that analysts need for ad-hoc sourcing decisions.

The emergence of large language models (LLMs) capable of structured reasoning, and the prolifera-tion of high-quality public and commercial APIs covering commodity markets, geopolitical events, and trade flows, creates an opportunity to construct a new class of procurement intelligence tool: one that is agentic in its data gathering, real-time in its signal processing, and accessible to procurement teams without enterprise IT infrastructure.

This paper makes the following contributions:

- We define and implement an agentic supply chain intelligence architecture that decomposes procurement intelligence into specialised API agents coordinated by a central reasoning layer.
- We describe a multi-source risk fusion methodology that combines geophysical, geopolitical, financial, and trade signals into a unified risk representation.
- We present NAUTILUS Terminal, a working implementation deployed at production scale, and evaluate its performance across 40 procurement scenarios.
- We discuss the design trade-offs involved in building real-time intelligence systems on heterogeneous public APIs, including reliability, latency, and data quality considerations.

Related Work

Supply Chain Risk Management Systems

Supply chain risk management (SCRM) has been extensively studied in the operations research literature [1-3]. Traditional frameworks categorise risk across supply-side (supplier failure, capacity constraints), demand-side (forecast error, demand volatility), and environmental (natural disasters, geopolitical) dimensions. However, the translation of these frameworks into operational software has historically been limited by data availability and integration complexity.

Recent work has explored data-driven approaches to supply chain risk. Brintrup et al. demon-strated graph neural network approaches to supply chain mapping from corporate filing data [4]. Ivanov and Dolgui proposed ripple effect models for disruption propagation across supply network tiers [5]. Neither approach addresses real-time signal integration or procurement decision support at the point of sourcing.

Agentic AI Systems

The concept of LLM-based agents, in which language models act as orchestrators of tool use and multi-step reasoning, has advanced rapidly since the introduction of ReAct [6, 7]. Multi-agent frameworks including LangChain and the Anthropic Agent SDK have demonstrated that compositional agent architectures can solve complex tasks requiring heterogeneous tool use. Applications of agentic AI to enterprise workflows have appeared in software engineering, customer support, and financial analysis. Supply chain applications remain nascent; Toorajipour et al. surveyed AI applications in supply chain management and identified a gap in real-time, LLM-mediated decision support systems [8].

Procurement Intelligence Tools

Commercial procurement intelligence tools may be broadly categorised as: spend analytics platforms (Coupa, Jaggaer), supplier risk monitors (Resilinc, Everstream), commodity price terminals (Bloomberg, Refinitiv), and trade compliance tools (Thomson Reuters ONESOURCE, Amber Road). These tools operate in isolation and require significant integration effort and licensing cost. NAUTILUS addresses this fragmentation through architectural unification of the

underlying data signals.

System Architecture

NAUTILUS Terminal is built on a layered agentic architecture comprising four strata: (1) a data in-gestion layer of specialised API agents, (2) a risk fusion and synthesis layer, (3) an LLM reasoning layer, and (4) a presentation layer. Figure 1 illustrates the high-level structure.

Table 1: NAUTILUS layered agentic architecture

Presentation Layer						
Next.js App Router	Tabbed Terminal UI	Interactive Globe Visualisation		PDF Export		
LLM Reasoning Layer						
Category Classification	Hub Recommendation	Directive Generation		UN Comtrade Enrichment		
Risk Fusion and Synthesis						
Geopolitical Score	Commodity Risk	Compliance Assessment		Environmental Alerts		
Commodity	**Geopolitical**	**Seismic**	**Wildfire**	**Maritime**	**FX**	**Intelligence**
Yahoo Finance / Stooq	ACLED	USGS	NASA FIRMS	26 Global Ports	Open FX	NewsAPI

Data Ingestion Agents

Each data domain is handled by a dedicated Next.js API route acting as an autonomous agent. Agents are stateless, independently cacheable, and expose a normalised JSON schema regardless of upstream API variation. Seven agents are currently deployed.

Commodity Agent (/api/commodities). Fetches live spot prices for ten commodities from Yahoo Finance, with automatic fallback to Stooq end-of-day data. Implements source quality classifica-tion (live, delayed, or reference) exposed as a quality field in the API response, enabling down-stream consumers to weight data appropriately.

Geopolitical Agent (/api/incidents). Integrates the Armed Conflict Location and Event Data (ACLED) API, retrieving geo-tagged political violence, protests, and conflict events within a con-figurable time window. Events are normalised to a four-level severity taxonomy (Critical, High, Medium, Low) based on event type and reported fatalities.

Seismic Agent (/api/earthquakes). Queries the USGS Earthquake Hazards API for seismic events above a configurable magnitude threshold. Severity mapping follows USGS ShakeMap intensity guidelines: M7.0 and above is Critical, M6.0 to 6.9 is High, M5.0 to 5.9 is Medium.

Wildfire Agent (/api/wildfires). Integrates NASA FIRMS near-real-time fire detection data. Active fire perimeters within proximity of known industrial zones are flagged as supply chain risk events with geographic coordinates.

Maritime Agent (/api/ports). Monitors congestion and operational status across 26 major global container ports, with risk scoring derived from historical throughput data and real-time incident feeds.

FX Agent (/api/fx). Fetches live foreign exchange rates for 15 major procurement currency pairs, with stale-data detection and fallback to last-known rates. FX volatility is surfaced as a procure-ment cost risk factor

within the compliance and cost panels.

Intelligence Agent (/api/intel). Aggregates supply chain-relevant news from NewsAPI, filtered through keyword extraction derived from the active procurement query. Implements abort-con-troller-based request management to prevent stale responses contaminating current results.

Risk Fusion and Synthesis Layer

Raw agent outputs are merged in the client-side risk fusion layer, which implements a multi-dimen-sional risk taxonomy with five categories: geopolitical, commodity, logistics, compliance, and envi-ronmental. Each risk entry carries a severity classification, a human-readable description, and a struc-tured mitigation recommendation. Critical events from any agent propagate to the terminal's primary alert surface regardless of query context.

LLM Reasoning Layer

The central intelligence function is implemented in the /api/analyze route, performing three distinct reasoning tasks.

Category Classification. The system maintains a lexical signal map of 40 procurement categories, each defined by domain-specific keywords. Free-text procurement queries are classified using to-ken overlap scoring, with confidence estimation to flag low-confidence matches for user review.

Hub Recommendation. The ATLAS_DB knowledge base encodes 40 sourcing hub profiles con-taining geographic location, key partner companies, logistics parameters, customs and HTS classi-fication data, ESG considerations, and risk descriptors. A weighted scoring function selects the op-timal sourcing region. For each recommended hub, the system queries the UN Comtrade API to at-tach real trade flow data, providing empirical grounding for recommendations.

Directive Generation. The system generates a structured sourcing directive comprising: the rec-ommended sourcing region, primary partner recommendation, logistics route specification, tariff alert with HTS code and duty rate, and a natural-language strategic summary. This directive is pro-duced deterministically from structured data, avoiding LLM hallucination risks.

Supporting Analytical Modules

Beyond the core agentic pipeline, NAUTILUS exposes eleven analytical modules: HS Code Lookup, Total Landed Cost Calculator, Supplier Comparison, Incoterms 2020 Guide, Compliance Checklist, Sanctions Checker, Ocean Freight Spot Rates, FTA Eligibility Checker, Tariff Calculator (with Section 301 and Section 232 surcharge overlays), FX Impact Calculator, and a Dual-Use Goods Checker for export control classification.

Implementation

Technology Stack

NAUTILUS Terminal is implemented as a Next.js 14 application using the App Router architecture. The choice of a unified full-stack JavaScript framework collocates API agent logic with the presenta-tion layer, eliminates cross-origin complexity, and enables edge deployment via Vercel's serverless in-frastructure. Table 1 summarises the technology stack.

Table 2: NAUTILUS technology stack

Component	Technology
Application Framework	Next.js 14 (App Router)
Rendering	React 18 with client/ server component split
Styling	Tailwind CSS
Globe Visualisation	Three.js (WebGL)
Charts	Recharts
Animation	Framer Motion
PDF Export	jsPDF
Deployment	Vercel Edge Network

Rate Limiting and Reliability

Each API route implements a sliding-window rate limiter (10 requests per 60 seconds per IP address) to prevent abuse and protect upstream API quotas. External API calls implement configurable timeout handling (default 8 seconds), graceful degradation to cached or reference data on failure, and struc-tured error responses that allow the frontend to surface data quality warnings without interrupting the user experience.

Data Quality and Transparency

A key design principle of NAUTILUS is data provenance transparency. Every live data panel displays a source freshness badge indicating the data origin, timestamp,

and quality tier (LIVE, EOD, or REF). This allows analysts to apply appropriate confidence weighting without requiring knowledge of the underlying API infrastructure.

Mission Persistence

Procurement scans are persisted as missions in browser local storage, enabling analysts to replay prior queries, compare results across time periods, and build an institutional archive of sourcing decisions. The mission data model captures query text, timestamp, primary hub recommendation, hub count, top partner, and the full sourcing directive. A terminal reset function clears the mission archive and re-stores the system to its initial state.

Evaluation

Classification Accuracy and System Latency

We evaluated NAUTILUS across 40 procurement queries spanning the full category taxonomy. Category classification achieved correct top-1 category assignment in 87.5% of queries. Low-confi-dence detection correctly flagged ambiguous queries in 91% of cases. End-to-end response latency averaged 2.3 seconds on Vercel's edge network, with a 95th-percentile latency of 4.1 seconds under concurrent load.

Case Study: Neodymium Magnets for EV Motor Assembly

To illustrate the system's operation, we trace a representative procurement query: neodymium magnets for EV motor assembly.

Classification. The query is assigned to the industrial category (primary signals: magnet, neodymium, NdFeB, motor) with high confidence.

Live risk context. The commodity agent reports rare earth prices up 6.8%. The geopolitical agent flags active incidents in Sub-Saharan Africa. The seismic agent reports no relevant events in pri-mary sourcing regions.

Hub recommendation. Baotou Inner Mongolia (China) is selected as the primary hub, based on an estimated 85% global NdFeB magnet production share. Secondary hubs: Japan (Shin-Etsu Chemical, TDK), Germany (VAC), Estonia (Neo Performance Materials).

Directive generated. Strategic scan complete for neodymium magnets for EV motor assembly. Identified 4 global sourcing hubs. Primary recommendation: Baotou Inner Mongolia. HTS: 8505.11.00. Duty: 3.7%. UFLPA supply chain audit required for all Chinese rare earth imports.

Trade flow enrichment. UN Comtrade data confirms China as the dominant exporter of rare earth compounds, corroborating hub selection with empirical trade volume data.

Comparison with Existing Tools

Table 3: Feature comparison across procurement intelligence tools

Capability	NAUTILUS	Bloomberg	Resilinc	SAP Ariba
Live commodity prices	Yes	Yes	No	No
Geopolitical event feed	Yes	Partial	Yes	No
Seismic and wildfire risk	Yes	No	Partial	No
LLM sourcing recommendation	Yes	No	No	No
Trade flow data	Yes	Partial	No	Partial
Compliance and sanctions	Yes	No	Partial	Yes
Landed cost modelling	Yes	No	No	Yes

No enterprise contract required	Yes	No	No	No
Response latency	<3 s	<1 s	Minutes	Minutes

Discussion

Agentic Design Trade-offs

The agentic architecture imposes meaningful trade-offs. Decomposing intelligence into independent API agents improves modularity and fault isolation: a failing wildfire API does not affect the com-modity or geopolitical agent. However, this increases the total external API call count per user ses-sion, raising quota consumption and aggregate latency relative to a monolithic data pipeline.

The decision to avoid LLM inference in the hot path of sourcing recommendations, using determinis-tic structured data instead, is a deliberate trade-off favouring reliability and cost over generative flexi-bility. LLM hallucination in a procurement context carries direct financial and compliance risk. The current architecture constrains LLM output to natural-language summarisation of deterministically generated structured data, substantially reducing this risk.

Limitations

Several limitations warrant acknowledgement. The ATLAS_DB sourcing hub knowledge base is man-ually curated and requires ongoing maintenance to remain current. The hub recommendation function does not incorporate dynamic capacity data, real-time supplier financial health signals, or Tier-2 and Tier-3 supplier mapping. Category classification relies on lexical matching rather than semantic em-bedding, which may produce incorrect classifications for highly specialised queries.

The system does not currently incorporate user authentication or multi-user state management. Server-backed mission persistence, OAuth-authenticated access, and collaborative annotation are identified as near-term development priorities.

Future Work

Semantic category classification. Replacing lexical keyword matching with dense retrieval over a procurement ontology would improve classification accuracy and enable cross-category risk inference.

Supplier graph integration. Incorporating Tier-2 and Tier-3 supplier dependency mapping using knowledge graph techniques would enable propagation of upstream disruption signals to down-stream procurement risk assessments.

Reinforcement learning from procurement outcomes. If procurement decisions and outcomes are logged, reinforcement learning from human feedback could be used to improve hub recom-mendation quality over time.

Federated deployment. A federated architecture in which organisations deploy private NAUTILUS instances with proprietary supplier data augmenting the shared public data layer could dramatically improve recommendation precision while preserving data confidentiality.

Conclusion

This paper presented NAUTILUS Terminal, an agentic supply chain intelligence architecture that ad-dresses the fragmentation of procurement data across geopolitical, geophysical, commodity, trade, and compliance domains. By decomposing intelligence gathering into specialised API agents, fusing their outputs through a structured risk taxonomy, and surfacing actionable directives through a real-time terminal interface, NAUTILUS demonstrates that production-grade supply chain intelligence is achievable without enterprise IT infrastructure

or multi-year implementation timelines.

The broader implication of this work is architectural: agentic composition of heterogeneous live data streams, mediated by structured LLM reasoning, represents a viable and practical approach to building real-time decision support systems in domains previously served only by expensive, siloed enterprise software.

The system is live and publicly accessible at https://nautilus-terminal.vercel.app. The ATLAS_DB knowledge base and system architecture are available for research collaboration upon request at anubhavtewari7@gmail.com.

References

1. Chopra S, Sodhi MS (2004) Managing risk to avoid supply-chain breakdown. MIT Sloan Management Review 46: 53-61.
2. Tang CS (2006) Perspectives in supply chain risk management. International Journal of Production Economics 103: 451-488.
3. Sodhi MS, Son BG, Tang CS (2012) Researchers' perspectives on supply chain risk management. Production and Operations Management 21: 1-13.
4. Brintrup A, Pak J, Ratiney D, Pearce T and Wichmann P, et al. (2020) Supply chain data analytics for predicting supplier disruptions: a case study in complex asset manufacturing. International Journal of Production Research 58: 3330-3341.
5. Ivanov D, Dolgui A (2021) OR-methods for coping with the ripple effect in supply chains during COVID-19 pandemic. International Journal of Production Economics 232: 107921.
6. Yao S, Zhao J, Yu D, Du N and Shafran I, et al. (2022) ReAct: Synergizing reasoning and acting in language models https://arxiv.org/abs/2210.03629.
7. Schick T, Dwivedi-Yu J, Dessi R, Raileanu R and Lomeli M, et al. (2023) Toolformer: Language models can teach themselves to use tools. Advances in Neural Information Processing Systems 36.
8. Toorajipour R, Sohrabpour V, Nazarpour A, Oghazi P and Fischl M (2021) Artificial intelligence in supply chain man-agement: A systematic literature review. Journal of Business Research 122: 502-517.

Copyright: ©2026 Anubhav Tewari. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.