DIGITAL SUPPLY CHAIN LAB

Tarion

AI Trade Strategy Copilot

KOSMIC 2026 | TOP 12 NATIONAL FINALIST

Project Development Report

A transparent sourcing and trade decision system for evaluating tariff exposure, landed cost, regulatory constraints, supplier risk, resilience, and feasibility.

Author	Shubham Gambhire
Affiliation	Indian Institute of Management Mumbai
Project period	8 July 2026 to 15 September 2026
Contact	shubhambgambhire@gmail.com
Portfolio	digitalsupplychainlab.com

Prepared for the Digital Supply Chain Lab portfolio. The competition certificate was used only to verify the author's recognition and is not reproduced in this report.

Contents

01 Executive Summary
02 Project Progression
03 Core Capabilities
04 Standards and Compliance
05 Issue Log and Fixes
06 Verification
07 Uses of the Tool
08 Honest Limitations

Executive Summary

Tarion is a browser-based trade strategy copilot that turns a tariff shock into a ranked, defended sourcing decision. It prices viable countries of origin for a product, screens options against thirteen hard rules, scores them on cost, risk, resilience, and feasibility, and keeps the arithmetic visible for selected and rejected options.

The project began as a case competition concept, then matured into a deterministic calculation engine with hand-verified regression tests. It expanded to two destination markets, two product families, live regulatory sensing with a human approval gate, supplier posture analysis, and a tightly controlled AI narration layer.

The work was completed as part of KOSMIC 2026, where Shubham Gambhire was recognised as a member of a Top 12 national finalist team.

The LLM explains, the engine decides. No mathematics is performed by a language model.

505	14	13	2
automated tests	hand-computed cases	hard rules	destination markets

Project Progression

Each stage preserved verified behavior while expanding scope. The report treats implementation evidence, not presentation claims, as the source of truth.

Stage 1: Lock the Logic

Specified a 21-step landed-cost calculation, thirteen threshold-based rules, a six-factor risk rubric, and a four-pillar ranking formula: Cost 40, Risk 25, Resilience 20, and Feasibility 15. A 30-SKU dataset used published tariff rates and real HTS codes. Fourteen cases were computed independently by hand before coding.

Stage 2: Build the Core

Built a pure Python engine with no web imports, a thin FastAPI layer, and a vanilla JavaScript interface. The separation is strict: the engine computes, the API transports, and the interface renders.

Stage 3: Generalise Product and Market Logic

Replaced hidden US-only assumptions. HS-6 became the portable product key, destination markets implement a DutyRegime interface, and EXW, FOB, CIF, and DDP quotes are normalised into one price stack with GIVEN, ADDED, or INFERRED provenance.

Stage 4: Meet Real Data Where It Is

Added row-level intake tiers, adjustable ranking weights, deterministic regulatory sensing with human approval, and synthetic fixtures that exercise previously dormant kill-check paths.

Stage 5: Add India and a Second Product Family

Implemented an India regime with ITC-HS codes, CIF valuation, and the BCD, SWS, and IGST cascade. Added twelve furniture SKUs across eleven HS-6 subheadings. Product assumptions became explicit data with QUOTED, ESTIMATED, and UNAVAILABLE states.

Stage 6: Add a Guarded AI Layer

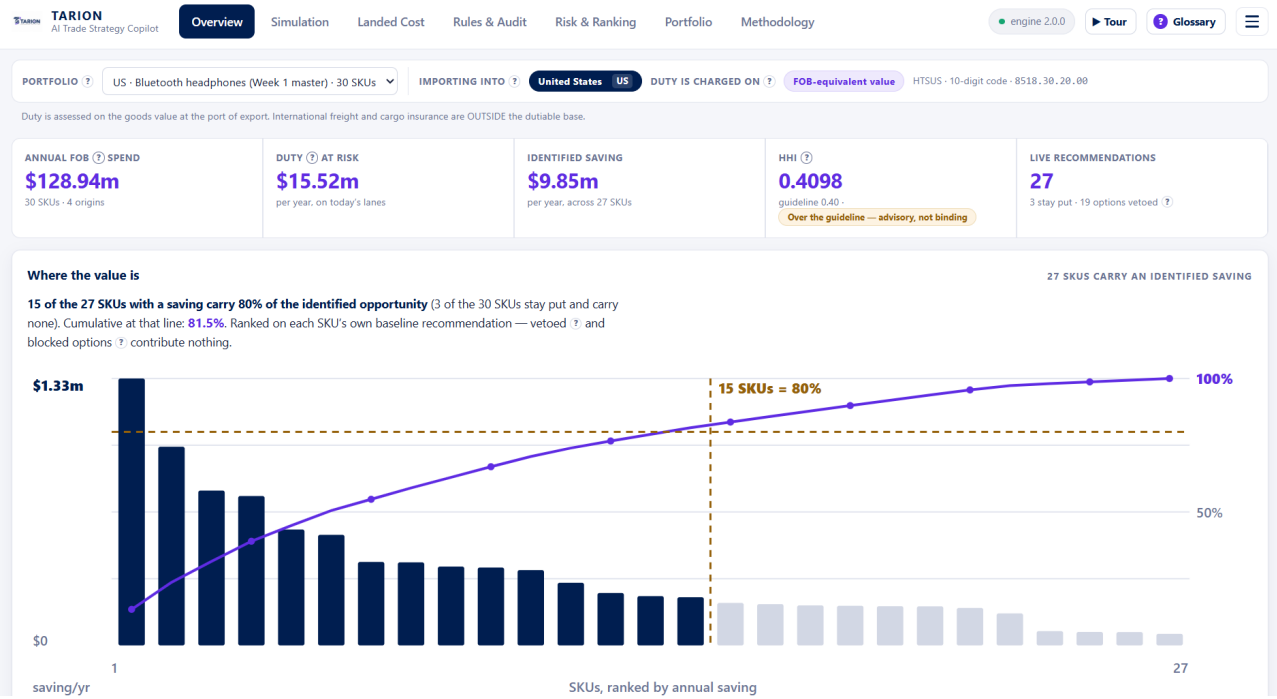
Opened only two model seams: regulatory text enters and a finished decision is narrated. The server recomputes every decision, then checks narration for invented numbers, re-ranking, and supplier endorsement. A failed narration never blocks the structured result.

Stage 7: Showcase, Audit, and Deploy

Reconciled 117 pitch claims against the build, kept only supportable language, assembled a 55-question challenge bank, and deployed the interface and Python engine as one public experience.

Product in Use

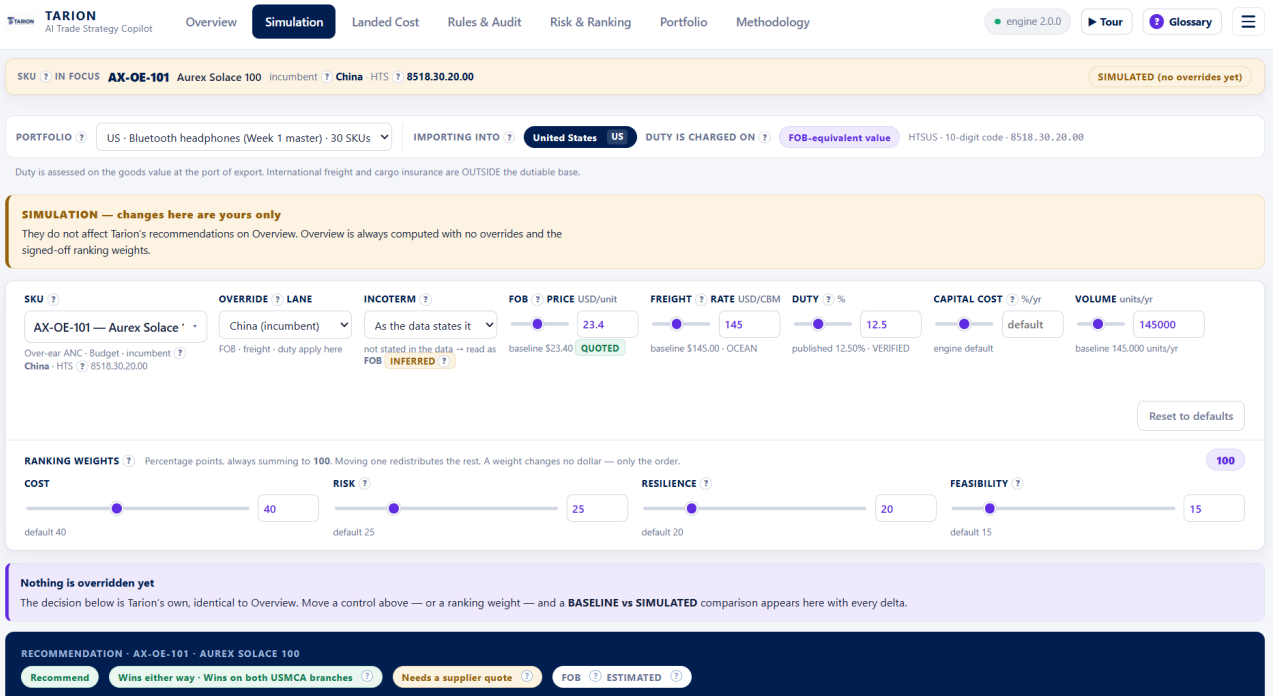
The overview concentrates attention on the portfolio value at risk, identified savings, concentration, and the few SKUs carrying most of the opportunity.



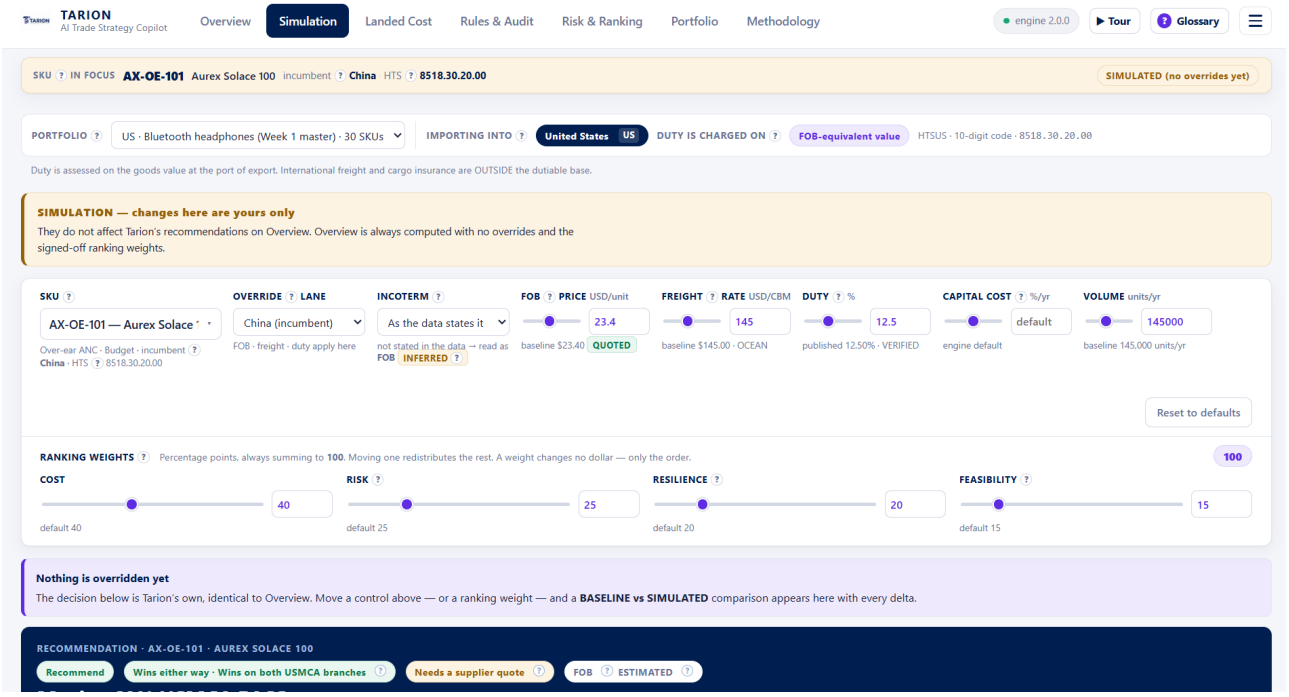
Portfolio overview with annual spend, duty exposure, identified savings, concentration, and a ranked savings curve.

Simulation and Decision Controls

Users can test lane, incoterm, price, freight, duty, capital cost, volume, and ranking-weight assumptions without changing the signed-off baseline recommendation.



Scenario controls preserve a visible separation between baseline and simulated states.



Recommendation evidence remains visible below the simulation controls.

Core Capabilities

Area	Capability
Landed cost	21-step cost build-up with full precision and output-only rounding.
Markets	US and India regimes with market-specific codes, valuation, fees, and preferences.
Incoterms	EXW, FOB, CIF, and DDP normalised to a nine-element stack with provenance.
Rules	Thirteen rules. Blocked and vetoed options remain visible with their arithmetic.
Risk and ranking	Six risk factors and four adjustable pillars with a documented tie band.
Data intake	T0 to T3 row-level quality assessment, fuzzy mapping, and capability reporting.
Regulatory sensing	Federal Register monitoring, deterministic parsing, and named human approval.
Supplier posture	Capacity, MOQ, qualification, lane concentration, and contract absorption.
AI narration	Server-side recomputation plus checks for invented numbers, re-ranking, and endorsement.
Interface	Eight screens, an 88-term glossary, and a twelve-step guided tour.

Standards and Compliance

Every rate, fee, and threshold is traced to a published source and carries a verification status. The model covers HTSUS and USITC classifications, Section 301 overlays, MPF and HMF, USMCA qualification, India customs duties and cascaded taxes, AIFTA preferences, UFLPA, and AD/CVD screening.

Rates are labelled VERIFIED, VERIFIED-SECONDARY, or UNVERIFIED-ASSUMPTION. The status travels with each value into the interface.

Issue Log and Fixes

Several findings changed the product conclusion rather than only improving code quality.

Finding	Resolution
Vietnam shift did not survive real rates	China and Vietnam were duty-identical. Rule R07 now vetoes the higher-cost lane and shows why.
India surcharge used the wrong base	Corrected SWS to 10% of Basic Customs Duty and pinned the result with a ground-truth case.
Transition risk was constant	Re-banded four components, producing 29 distinct values and changing rankings.
Unsatisfiable concentration rule	Reclassified as advisory while retaining the binding incumbent floor.
Weights existed in three places	Moved to one source so ranking and explanation cannot diverge.
Simulation controls were not wired	Connected initialization and added browser-level verification.
Non-headphone products raised an error	Replaced product-keyed dictionaries with explicit provenance-based data.
All supplier prices assumed FOB	Added incoterm normalisation to prevent silent mixed-quote errors.
Internal files risked public deployment	Added a public-root audit that permits only known web assets.

Verification

- 505 automated tests with zero skipped.
- Fourteen hand-computed ground-truth cases reproduce to four decimal places after every change.
- A separate eight-case oracle validates the India regime without importing the engine.
- The signed-off dataset is pinned by SHA-256.
- A browser-level smoke test renders the interface against real API responses and fails on uncaught exceptions.

Uses of the Tool

- Tariff shock response by SKU.
- Sourcing strategy review across landed cost, risk, resilience, and feasibility.
- Procurement negotiation support with line-item input provenance.
- Regulatory monitoring with a human-reviewed audit trail.
- Portfolio concentration analysis by country and feasible lane capacity.
- Interactive teaching of trade and sourcing economics.

Honest Limitations

- The portfolio dataset is synthetic, although tariff rates, HTS codes, and fee constants use published sources. Savings are illustrative, not realised.
- No live LLM is attached. The AI integration seams and guardrails are demonstrated through deterministic implementations.
- Some India rates are secondary-source verified, and one preferential rate remains an explicit unverified assumption.
- Freight, transit, and compliance constants are calibrated for US-bound lanes, so absolute India landed costs are indicative.
- Supplier capacity, MOQ, and qualification attributes are derived in the demonstration dataset.
- The tool has not been used by a real customer and has no production usage or retention evidence.

Conclusion

Tarion demonstrates a disciplined approach to AI-enabled decision support: transparent calculations, explicit assumptions, visible rejection logic, human approval for regulatory changes, and language models restricted to explanation. Its strongest contribution is not a single recommendation, but a defensible process for comparing sourcing choices under trade uncertainty.