

Product Requirements Document – Risk Radar

1. Product Overview

Objective

A real-time dashboard that tracks global supply chain disruptions, scores their severity, and visualizes them on a map. The goal is to give managers a single, clear view of what matters, fast.

Primary Users

- Supply chain managers planning logistics.
- Risk analysts monitoring disruptions.
- Business leaders making trade or sourcing decisions.

Problem Statement

Right now, disruptions are scattered across dozens of sources. There's no consistent way to judge severity. Decision-makers waste hours piecing together fragmented updates — and still risk missing what matters.

2. Goals & Non-Goals

In Scope (MVP)

- Ingest news from public APIs.
- Classify events into disruption types.
- Score events with a Supply Chain Severity Index (SCSI).
- Visualize on a map with filters for region, type, severity.

Out of Scope (for now)

- Predictive modeling of future disruptions.
- Integration with private trade or ERP databases.
- Automated sentiment analysis.

Why: Kept the scope lean to deliver real-time awareness quickly.

3. User Stories

- *As a supply chain manager, I want a global map of disruptions ranked by severity, so I know which events demand immediate response.*
- *As a risk analyst, I want consistent scoring rules, so I can compare disruptions objectively instead of arguing about severity.*
- *As a decision-maker, I want filters by geography and type, so I can focus only on disruptions relevant to my operations.*

4. Solution Overview

Data Flow

- NewsAPI fetches disruption-related news.
- n8n pipeline cleans and deduplicates articles.

- LLM (Gemini) classifies disruption type, impacted nodes, severity sub-scores.
- Google Sheets applies scoring rubric (0–10 SCSI with decay and caps).
- Looker Studio connects live data for visualization.

Severity Index

Weighted model with five factors:

- Scope (25%)
- Node Importance (25%)
- Risk Type (20%)
- Trade Dependency (20%)
- Population Impact (10%)

Adjustments: age decay, confidence caps, speculative caps, and spread bonus.

Interactivity

- Map view with color-coded severity.
- Ranked table of top events.
- Trend line of disruptions over time.
- Drill-downs with source link and severity breakdown.

5. Success Metrics

North-Star: Time to risk identification (reduce from hours/days to near real-time).

Drivers:

- $\geq 80\%$ accuracy in event classification.
- Coverage of all major disruptions within 24 hrs.
- Dashboard refresh latency < 5 mins.

Counters:

- False positives kept under 10%.
- User drop-off (session time < 1 min).

6. Discovery Insights

Early user conversations highlighted:

- *“We track risks on WhatsApp groups and Google Alerts — it’s hit or miss.”*
- *“Severity is always subjective. Everyone has their own scale.”*
- *“Even if I know something happened, I don’t know if it matters to my lane.”*

These insights drove the focus on a transparent severity index and simple, interactive map.

7. Trade-Offs & Risks

Trade-Offs

- Chose rule-based severity index over predictive ML for transparency.
- Opted for public data feeds (NewsAPI) over paid/private feeds to keep it low-cost and replicable.
- Prioritized classification accuracy over alert automation in v1.

Risks

- Dependence on news coverage → some events may be under-reported.
- LLM misclassification risk, mitigated by conservative scoring caps.
- Data freshness limited by API refresh rate.

8. Roadmap

- **v1 (MVP – Done):** Live dashboard with ingestion, classification, scoring, visualization.
- **v2:** Add multi-language feeds, basic alert system (email/Slack).
- **v3:** Predictive modeling, integration with private trade data, advanced analytics.

9. Appendix

- Full severity scoring framework.
- LLM prompt design.
- Sample JSON output.