

One simple question activates one of the world's most advanced geopolitical simulation engines.

A unified intelligence platform that measures the state of the world in real time across security, cyber, health, finance, politics and narrative — and turns the answer into a decision.

THE PROBLEM

Collection is solved. Processing is not.

01 The volume outgrew the analyst

A mid-sized agency now takes in more open-source signal in a day than a Cold War station chief saw in a year. Collection scaled. The human capacity to weigh it did not.

02 Analysis that cannot be wrong

"Tensions may rise in the coming period" cannot be falsified, so it cannot be trusted. Almost no institution keeps a ledger of what its analysts predicted and what actually happened.

03 Prose, not a decision

The output arrives as twenty pages. The reader still has to convert it into an action by Thursday, under time pressure, alone, and carry the accountability for it.

Everyone has collection. Processing is the difference.

The bottleneck moved years ago. Most of the market is still paying for the old one.

THE SOLUTION

One question in. Three weighted futures out. Then the tasks.

ASK

Plain language

No syntax, no training, no analyst in between. You ask the way you would ask a colleague.

BRIEF

~45 seconds

A full read across nine domains and six lenses, including competing ideological perspectives.

SIMULATE

~4 minutes

Three forward scenarios with early-warning indicators, lead times, falsifiers and a stress index.

ACT

The workspace

Tasks, playbook, associated risks and, where tempo demands it, a 24-hour plan.

The model never proposes a number. Probabilities are computed from base-rate priors updated by weighted signals, log-odds pooled and calibrated against observed outcomes — the language layer presents them, it does not invent them. Every output carries HMAC/SHA-256 provenance, and users can inject their own intelligence, source-graded to NATO standard and classified. Their edge stays theirs.

COVERAGE

NINE DOMAINS

Security · Political · Economic · Climate · Narrative · Cyber · Health · Stability · Space

SIX LENSES

Political realism · Constructivism · Honour-shame · Clan dynamics · Cognitive warfare · Decision psychology

WHAT IT DOES

Generate simulations

Test geopolitical and security outcomes before they happen. Decision paths built from real-time signals, early-warning indicators and probabilistic outcomes — foresight rather than hindsight.

Global Stress Index

One score, one continuous signal. GSI quantifies global and regional stress across security, cyber, health, finance, politics and narrative pressure, in real time and rescored continuously.

Site monitoring

Automatic monitoring of what happens around each site or office: threat assessments, live headcounts and orchestrated emergency response without manual intervention.

Travel

Journey-specific simulations showing how evolving global conditions will affect a trip before it begins, at the level of route, city and date. The system warns while something is building, not once it has happened.

Workspace

A digital head of security and project lead — turning questions and signals into decisions, workflows and executed tasks, with recommendations one click from the knowledge base.

50+

tailored AI agents working simultaneously

10M+

signals collected and continuously scored

72h+

of early warning before escalation

~4 min

to generate a full simulation

GSI

one live index, rescored continuously

10h+


saved per decision-maker per week

90%+

of users report higher situational readiness

3

scenarios distilled from thousands of futures



The working surface: one plain-language question, nine domains, six lenses, three weighted futures.

WHY US

No prompt replicates this. It is not a prompt — it is a career.

GENERIC LLM / TRADITIONAL PROVIDER

One pass. No memory of truth.

- No persistent signal ingestion — one pass over training data and a single search.
- First answer is the only answer. No adversarial self-challenge.
- Confident prose, unverifiable provenance, no source-tier discipline.
- Raw data or a written report. The reader still does the analysis.
- Built by people who never had to act on their own assessment.

OMEGANOVA

Built by the operator, not only the engineer.

- 25 years of international security work and more than 4,000 missions in high-risk environments, encoded as domain logic.
- Operational manuals, asymmetric doctrines, surveillance patterns and decision psychology — curated, not scraped.
- Probabilities computed, calibrated and scored against outcomes. Misses are published alongside hits.
- Nine domains, six lenses. A clan succession read through Western institutional logic is not a hard problem — it is the wrong one.
- Sits on top of existing situational awareness. No integration into classified networks required to prove value.

Palantir makes your data analysable. OmegaNova makes the external world assessable.

Different axis — not a smaller version of the same thing. We sit on top of what you already own.

PROOF

— Nepal, 9 September 2025

The escalation path was flagged 47 hours before the parliament building in Kathmandu burned. Forty-seven hours is an evacuation, a moved delegation, a warned staff.

— Iran

A full escalation chain called in sequence, with an 85-day lead on the opening move.

Every decidable case is logged in a public register — hit, partial, miss or void — with the misses visible. Aggregate accuracy figures publish after the register freeze. Buyers to date span a civil government agency, a defence engagement and a multilateral tender: three different buyer types, one engine.



Name a scenario. We simulate it live, in four minutes.

Any question, any region, in the room. No preparation, no dataset required from you.

omeganova.ai