

AI You Can Defend to a Regulator

Your GRC platform produced a risk score this morning. A regulator asks how. That question is where most AI in governance, risk, and compliance stops being an advantage and starts being an exposure.

The Mitratesh GRC AI Ecosystem



OPT-IN AT THE FEATURE LEVEL

Toggle AI capability by function with Mitratesh solutions. Turn it on, turn it off, change your mind next quarter. Nothing runs because a vendor decided it should.



TRACEABLE TO SOURCE

Mitratesh ARIES™ works from structured, tested APIs against your real GRC data. Every response traces back to source data your team can verify.



HUMAN-IN-THE-LOOP BY DESIGN

AI populates the assessment. Your team sees exactly which answers it wrote, which questions it could not, and accepts or rejects each one.



YOUR DATA STAYS YOURS

No client data is used to train Mitratesh AI models. Access runs through SSO with role-based permissions, so answers respect what each user may see.

Five AI Questions to Ask Every GRC Vendor

1

Can your system create a tamper-proof audit log for every AI decision, showing the exact system state, intent, and tools used?

2

Are your safety guardrails built into the core system infrastructure, or do they rely on unpredictable AI model behavior?

3

Can you show a real-world example of how your platform traces an AI output back through its step-by-step reasoning?

4

How do you track prompt and model versions so an AI decision made six months ago can be recreated during an audit?

5

Does your contract protect us against regulatory fines caused by automated AI actions, or do you limit your liability?

Where AI Works Across Your GRC Program

Mitratech ARIES™ unifies data across 11+ domains with the Global GRC Platform, giving you one answer in seconds instead of logging into each system separately. It also powers the AI inside every product below, from Mitratech Alyne's Risk Agent to Mitratech Prevalent's AI False Positive Detection. Here is where it does the work.

ENTERPRISE RISK AND ASSESSMENTS

Mitratech Alyne

Mitratech Alyne maps your policies and evidence to the controls and questionnaire items they touch, then recommends an answer for each open item grounded in what you uploaded. Your team reviews every recommendation before it counts. Evidence analysis cuts completion time by 25 to 40 percent on long-form questionnaires, and exposure gets quantified through a built-in simulation engine.

CROSS-PRODUCT ASSESSMENT REUSE

Mitratech Global GRC Platform

Complete an assessment in one product. Its answers carry into a new assessment anywhere else on the platform, no matter where it started. That only works because the Mitratech Global GRC Platform connects the data underneath every product into one layer, instead of leaving each one siloed.

BOARD REPORTING

Mitratech ARIES™

Mitratech ARIES™'s Risk Agent builds executive reports directly from your risk library and answers ad-hoc board questions in minutes. Every agent action stays scoped to your permissions and gated by explicit human approval before it runs.

THIRD-PARTY RISK

Mitratech Prevalent

AI extracts key clauses from vendor documents, scored for confidence with a page reference you can check. Vendor surveys get a first-pass response grounded in evidence the vendor uploads, for your team to confirm. Administrators can turn any AI feature on or off at any time.

AI RISK IN YOUR RISK REGISTER

Mitratech Alyne + Mitratech ClusterSeven

Find the AI and ML models running in your organization and across your third parties, including the ones outside IT's control, alongside the spreadsheets and end-user computing your reporting depends on. Score them in the same register as every other enterprise risk and assess against the NIST AI Risk Management Framework. Board-level visibility of where AI creates exposure.

POLICY, ETHICS, AND RESILIENCE

Mitratech PolicyHub · Mitratech Syntrio · Mitratech Prepara

Policy, ethics, and resilience share one data intelligence layer instead of running as three separate tools. Policy documents get AI summaries and better search, so reviewers find what they need without opening the whole file. Hotline reports get the same summary treatment, with key details highlighted and categories flagged automatically. Continuity plans get checked against standard frameworks on their own, cutting a review that took hours down to minutes.