

CLAIR™

Cyber Liability AI Risk Advisor

An AI companion that sits inside the cyber insurance conversation: translating policy language, pressure-testing controls, and turning a client discussion into a submission an underwriter can actually read.

Developed by **Avanew**

Deployment **Hosted** · **Embedded** · **White-label**

What CLAIR is

CLAIR is a conversational interface to the cyber insurance market. Not a quoting engine, and not a chatbot bolted onto a marketing site. A user asks in plain language; CLAIR answers in the language of the policy, cites the clause or data source behind the answer, and keeps the thread so the next question builds on the last.

Retail brokers and producers

Answer tier-one client questions instantly; arrive at the renewal meeting with the coverage conversation already half done.

Underwriting and wholesale teams

Receive submissions that are complete, consistent, and framed against the controls that actually move price.

Insureds, CISOs and risk managers

See which controls actually move price, where the gap sits, and what evidence a carrier will want, before the application is opened.

Core capabilities

■ Policy language translation

Explains forms, endorsements, sublimits and exclusions, from war and widespread event to betterment and prior acts, clause by clause.

■ Coverage-part explainers

Network security and privacy liability, business interruption and contingent BI, extortion, funds transfer fraud, media, regulatory defense.

■ Control readiness walkthrough

Steps a client through MFA scope, EDR, immutable backups, privileged access, email filtering, patch cadence and IR testing.

■ Gap flagging before submission

Identifies the answers most likely to trigger a referral, a subjectivity, or a declination.

■ Exposure narrative

Builds a scenario story scaled to industry, revenue band and data footprint, with modeled frequency and severity ranges.

■ Limits and retention framing

Supports the "how much is enough" discussion with modeled loss exceedance rather than peer benchmarks alone.

■ Claims-path navigation

First-72-hours sequence, notice provisions, panel vendor selection, consent-to-settle mechanics.

■ Session-to-submission output

Converts the conversation into a structured, citable summary the placement team can attach to a submission.

What a user can ask it

"My carrier added a widespread event exclusion. What does that actually take away?"

"We run MFA on email but not on remote access. How much does that hurt us?"

"What is the difference between contingent BI and system failure coverage?"

"We are a 900-employee manufacturer at \$400M revenue. Is a \$10M limit defensible?"

"Walk me through what happens in the first 24 hours of a ransomware notice."

How a session works

Frame

CLAIR opens by establishing sector, size, data types and current program structure. Three or four questions, not a form.

Converse

The user asks freely. Every answer carries an inline citation to a form, clause, threat source or modeled figure.

Pressure-test

CLAIR surfaces the gaps the conversation exposed and states plainly which ones a carrier will price or refer.

Hand off

The session exports as a structured summary for the broker or underwriter: questions asked, gaps found, coverage points raised.

Under the hood

REASONING LAYER

Large language model constrained by retrieval. CLAIR answers only from its curated corpus and says so when the corpus is silent.

KNOWLEDGE CORPUS

Standard and manuscript policy forms, carrier appetite guides, regulatory notification regimes, published claims precedent, threat intelligence and industry loss data.

QUANTIFICATION

Frequency and severity distributions expressed as loss exceedance ranges, refreshed on a scheduled cadence.

GUARDRAILS

Deterministic refusal on quoting, binding, and legal advice. Every substantive answer is citation-linked; unsupported answers are withheld rather than approximated.

INTERFACE

Browser-based chat, embeddable in a broker portal or client extranet. No installation. Full session transcript available to the user.

SECURITY AND PRIVACY

- **No training on customer data.** Session content is never used to retrain the underlying model.
- **Encryption** in transit and at rest, with tenant-level key separation.
- **Configurable retention**, from ephemeral sessions to a retained audit trail, set per deployment.
- **Regional data residency** available per deployment.
- **Role-scoped access** so a client-facing instance cannot surface internal appetite or pricing material.

Attestation status and deployment specifications confirmed by Avanew per engagement.

WHAT CLAIR DOES NOT DO

- It does not **quote, bind, or issue** coverage.
- It does not give **legal advice** or render a coverage determination on a live claim.
- It does not **replace the broker**. It removes the repetitive half of the conversation so the judgment half gets more time.
- It does not **guess**. Where the corpus has no answer, it says so and routes the question to a human.

VALUE PROPOSITION

Every client conversation starts from an informed client.

The cost of cyber insurance is not only premium. It is the hours spent explaining the same eight concepts, the submissions returned incomplete, and the disputes that trace back to a clause nobody read aloud. CLAIR moves that knowledge to the front of the process, where it is cheap, instead of the back, where it is expensive.

Talk to CLAIR now → [Avanew](#)

Private walkthrough → jeffrey.sharer@sharerconsulting.com