JIEGOU

JIEGOU OPERATOR BRIEF

Cyber Underwriting Readiness for AI

An Operator's Brief for 2026 Renewals

VERSION	v1.2 · 2026-05-23
FROM	JieGou — AI Operations Partner for engineering-led mid-market IT teams ($50M–$1B revenue)
FOR	CIOs · CTOs · CISOs · CFOs · Risk leaders preparing for 2026 cyber insurance renewal cycles
STATUS	Public artifact — free to use, cite, and circulate

jiegou.ai

NO SALES-CALL CONDITION

What this brief is

The AI section is now in your cyber insurance underwriting questionnaire. Whether you're aware of it yet depends on your renewal cycle and your broker. This brief maps the questions you'll be asked to the operator-grade answers you should be giving — anchored on the 10-Layer AI Governance framework JieGou uses internally and with every customer.

This isn't a sales brochure. It's the document we'd want our own CIO to have if we were on the buyer side of this conversation. It's structured so you can hand it to your broker as a starting point for your AI-section submission, OR run it yourself as an internal readiness check before your next renewal call.

Time to read: 20-30 minutes. Time to use: 1-3 hours, depending on how much of your AI estate is already documented.

---

§1 — Why this brief exists now

The structural change that made this brief necessary happened between mid-2024 and early 2026: every major broker firm publishing market commentary now explicitly identifies AI as a 2025-2026 underwriting focus area. Four named citations anchor that claim:

> ***"Underwriting reviews are now sharply focused on control maturity, vendor dependencies, AI use, and privacy practices."*** *— Aon, 2026 cyber and E&O market commentary*

> ***"Underwriters are focusing more on AI exposures and have had to adapt how they underwrite to AI exposure, trying to better understand how insureds utilize AI by asking a broader range of questions."*** *— Marsh, Q4 2024 US Cyber Market Update*

> ***"Underwriters are scrutinizing board and senior management oversight of AI governance. Insurers are not only asking questions about documented policies regarding AI usage but also innovating around AI and clarifying policy language."*** *— Lockton, December 2025 Cyber Market Update*

> *Underwriting questions have evolved from "Do you use AI?" to* ***"What models are you currently utilizing? How did the business decisions get made to utilize these models? What checks and balances are in place to make sure the AI models produce accurate and verifiable results?"*** *— Insurance Business, 2026 reporting on AI risk in cyber insurance*

In parallel, regulators have moved. NYDFS issued an October 2024 Industry Letter explicitly addressing "Cybersecurity Risks Arising from Artificial Intelligence" — requiring Senior Governing Body oversight of AI cybersecurity programs. The NAIC Model Bulletin on AI Use by Insurers has been adopted by 24 states as of 2026, requiring documented governance structures. The EU AI Act creates downstream pressure for any organization with European operations.

The translation: regulatory pressure on **insurers** becomes underwriting-question pressure on **insureds.** Your underwriter has new boxes to check; your submission has to answer them or risk being placed on a default exclusion tier.

This brief is operator-honest about what we can claim and what we can't. **We do not claim our framework reduces your premium by X%** — no insurer publicly offers vendor-specific governance discounts as of May 2026. What we do claim, and what this brief evidences: documented AI governance maturity is increasingly being asked about, and a structured framework outperforms an ad-hoc memo when a broker is assembling your submission packet.

---

§2 — The six AI underwriting question categories appearing in 2026 mid-market submissions

The actual mid-market cyber submission questionnaires are broker-confidential, but the question shape is openly discussed in broker market updates. The six categories below are synthesized from Aon, Marsh, Lockton, WTW, and Coalition public materials, plus the NYDFS Industry Letter framework.

Category 1 — AI Use Inventory and Acceptable-Use Policy

What underwriters ask:

- Does your organization use generative AI in production workflows? If yes, describe.
- Do you have a written AI acceptable-use policy?
- Are employees trained on the policy? When was the most recent training cycle?
- What's your enforcement mechanism for shadow AI (unauthorized AI tool use)?

Why this matters: "Shadow AI exclusion language" has emerged in 2025-2026 — unauthorized AI use can be treated as gross negligence if no written policy exists. The presence of a policy is increasingly a

coverage-eligibility floor, not a nice-to-have.

Which 10-Layer Governance layers address this: Layer 7 (Compliance) — documented policies, training records, enforcement mechanism. Layer 1 (Identity & Access) — SSO/SAML restricting AI tool access enforces the policy technically.

Operator-grade answer shape: "We maintain a written AI acceptable-use policy version-controlled in our compliance system, with documented quarterly training cycles for all employees with AI tool access. Shadow AI exposure is technically constrained via SSO restriction of AI tools to managed accounts."

Category 2 — AI Governance Framework and Senior Oversight

What underwriters ask:

- What governance framework do you use for AI risk? (NIST AI RMF, EU AI Act, ISO 42001, or other?)
- Have you conducted an AI risk assessment in the last 12 months? Provide documentation.
- Who in your organization owns AI governance at the senior level?
- What's the cadence of board / executive review of AI governance posture?

Why this matters: Lockton (Dec 2025) specifically calls out that underwriters are "scrutinizing board and senior management oversight of AI governance." This is structural — carriers now expect documented oversight chains analogous to financial controls.

Which 10-Layer Governance layers address this: Layer 7 (Compliance) — framework adoption + mapping. Layer 9 (Observability) — measurement / reporting cadence. Layer 10 (Incident Response) — incident review process.

Operator-grade answer shape: "We use the 10-Layer AI Governance framework, which cross-maps to SOC 2 Common Criteria (CC6 + CC7), EU AI Act risk tiers, NIST AI RMF (govern / map / measure / manage), and ISO/IEC 42001. We complete a 30-question self-assessment quarterly with results reported to executive leadership. Most recent assessment dated [DATE]."

Category 3 — Access Controls and Per-Agent Identity

What underwriters ask:

- Are AI tools integrated with your SSO/IAM?
- Do AI agents have separate identities from human users?
- What's your access-revocation process when an AI tool is removed from production?
- Is there role-based access control (RBAC) governing AI tool usage?

Why this matters: Per-agent identity is the part most platforms get wrong. If "the AI did it" collapses into a single audit subject, you can't answer the next question — and that's the question that drives the breach forensics + insurance claim narrative.

Which 10-Layer Governance layers address this: Layer 1 (Identity & Access) — SSO/SAML for AI consoles, per-agent identity, RBAC for AI agent management. Layer 6 (Tool Governance) — access controls on which tools agents can invoke.

Operator-grade answer shape: "AI tools integrate with our enterprise SSO/SAML. Each AI agent has its own scoped identity distinct from the user who created it, with per-tool permission sets. Agent-identity provisioning + revocation flows through our identity governance platform with full audit trail."

Category 4 — Audit Trail and Evidence Emission

What underwriters ask:

- What logging exists for AI agent actions?
- Can you produce an audit trail for any AI-driven decision in production?
- What's your audit-trail retention policy?
- Can you export evidence to your SIEM (Splunk / Sentinel / syslog)?

Why this matters: Post-breach, the insurance claims process turns on the question "what did the AI do, when, on whose authority, with what data?" If you can't reconstruct the chain, claim adjudication becomes adversarial. Underwriters increasingly underwrite the answer to this question, not just the existence of logging.

Which 10-Layer Governance layers address this: Layer 2 (Audit Trail) — full attribution, traceability to source data + prompt + model version, evidence export. Layer 9 (Observability) — real-time monitoring + dashboards.

Operator-grade answer shape: "Every AI agent action is logged with full attribution: actor identity (human or per-agent), input data, prompt template version, model version, output, timestamp. Logs are HMAC-signed (hash-chain integrity) for SOX/FDA/EU AI Act evidentiary contexts. Append-only by design; exported to our SIEM in real time."

Category 5 — Vendor Risk and LLM Provider Register

What underwriters ask:

- Do you maintain a vendor risk register for your AI providers (LLM, MLOps, etc.)?
- What's your data handling agreement with each AI vendor?

- Is data processed by AI vendors stored in approved residency regions?
- What's your exit strategy if a primary AI vendor becomes unavailable?

Why this matters: Lloyd's of London (2025) explicitly identified concentration risk: "high concentration of insurers using the same cloud providers and similar Large Language Models" as an operational resilience concern. The corollary at the insured level: carriers now ask whether you have provider diversity + an exit plan.

Which 10-Layer Governance layers address this: Layer 5 (Model Governance) — multi-provider evaluation, provider-portable workflows. Layer 3 (Data Governance) — data residency controls, PII/PHI detection. Layer 10 (Incident Response) — vendor risk register, sub-processor list.

Operator-grade answer shape: "We maintain a vendor risk register covering all AI providers with documented data-handling agreements, residency commitments, uptime SLAs, and sub-processor lists. Production workflows are designed to be model-agnostic — switching primary LLM provider does not require workflow rewrites. We can demonstrate this with an active multi-provider deployment."

Category 6 — Incident Response and Exclusion-Language Risk

What underwriters ask:

- Do you have an incident response plan covering AI-specific failures?
- What's your dead-letter / retry mechanism for failed AI operations?
- Have you tabletop-tested an AI-related incident scenario in the last 12 months?
- Are AI incidents tracked separately from general security incidents?

Why this matters: The carrier-side stick. AI exclusions are now showing up in commercial liability (ISO CGL forms CG 40 47 / CG 40 48 / CG 35 08 went live January 2026 with 80%+ state regulator approval; Chubb / Travelers / Berkshire reportedly filing similar in cyber). The cleanest defense against AI exclusion risk is documented incident response specific to AI failures, plus tabletop evidence.

Which 10-Layer Governance layers address this: Layer 10 (Incident Response) — DLQ, incident tracking, vendor risk register. Layer 4 (Human Oversight) — approval gates + escalation policies. Layer 8 (Cost Controls) — rate limiting prevents runaway-loop cost incidents.

Operator-grade answer shape: "Our AI-specific incident response plan covers prompt injection, data exfiltration, model jailbreak, audit-trail tampering, cross-tenant leakage. AI incidents are tracked separately from general security incidents with named accountable owners. Most recent tabletop exercise [DATE] simulated [SCENARIO]."

---

§3 — Mapping the 10 Layers to the 6 Question Categories

Layer	Primary categories addressed	Secondary categories
1 — Identity & Access	Cat 3 (Access controls)	Cat 1 (Acceptable use enforcement)
2 — Audit Trail	Cat 4 (Evidence emission)	Cat 6 (Incident reconstruction)
3 — Data Governance	Cat 5 (Data residency in vendor agreements)	Cat 6 (PHI/PII in incident scope)
4 — Human Oversight	Cat 6 (Approval gates as incident prevention)	Cat 2 (Senior oversight chain)
5 — Model Governance	Cat 5 (Vendor risk + provider portability)	Cat 2 (Framework adoption)
6 — Tool Governance	Cat 3 (Tool-level access controls)	Cat 1 (Shadow AI prevention)
7 — Compliance	Cat 1 (Policy), Cat 2 (Framework + assessment)	All categories (cross-mapping)
8 — Cost Controls	Cat 6 (Runaway-loop incident prevention)	—
9 — Observability	Cat 4 (Real-time monitoring)	Cat 2 (Reporting cadence)
10 — Incident Response	Cat 6 (Incident plan + DLQ + register)	Cat 5 (Vendor risk register)

The pattern: **every underwriting category maps to multiple layers; every layer addresses at least one category.** This is the structural argument for why a single framework outperforms category-specific point solutions when assembling an underwriting submission.

---

§3.5 — The AI Exclusion Landscape (2026 filing-verification update)

Underwriters asking about AI is one half of the story. The other half: carriers actively filing exclusion endorsements that strip AI exposure out of coverage at renewal. This section documents what's actually been filed with state regulators — verbatim, with form numbers and source links — and separates verified primary-source filings from trade-press claims that are circulating without independent verification.

Two distinct exclusion styles in the market

The exclusions in active state filings as of May 2026 fall into two structurally different categories. Conflating them in your renewal preparation will get you wrong answers in two directions at once.

Style A — ISO narrow exclusion (Commercial General Liability line, GenAI only)

The Insurance Services Office (ISO) — whose forms hundreds of US carriers license — published three new commercial general liability endorsements effective **January 1, 2026** (multistate). All three share an identical narrow definition of generative AI; they differ only in which CGL coverage they exclude:

ISO form	Coverage excluded	Scope
CG 40 47 01 26	Coverage A (BI/PD) + Coverage B (P&AI)	Broadest CGL exclusion; both coverages
CG 40 48 01 26	Coverage B (Personal & Advertising Injury) only	Defamation, copyright/trademark in ads, etc.
CG 35 08 01 26	Products / Completed Operations BI or PD	Downstream defect / failure-to-warn claims from GenAI in delivered products

Shared ISO definition (identical verbatim text across all three forms):

> *"'Generative artificial intelligence' means a machine-based learning system or model that is trained on data with the ability to create content or responses, including but not limited to text, images, audio, video or code."*

Verbatim operative exclusion text — CG 40 47 (the broadest):

> *"A. The following exclusion is added to Paragraph 2. Exclusions of Section I – Coverage A – Bodily Injury And Property Damage Liability: This insurance does not apply to: 'Bodily injury' or 'property damage' arising out of 'generative artificial intelligence'.*
>
> *B. The following exclusion is added to Paragraph 2. Exclusions of Section I – Coverage B – Personal And Advertising Injury Liability: This insurance does not apply to: 'Personal and advertising injury' arising out of 'generative artificial intelligence'."*

Source: ISO form CG 40 47 01 26, verbatim via PropertyCasualty360 / FC&S Bulletins (the same publisher that hosts the CG 40 48 and CG 35 08 SAMPLE PDFs). CG 40 47 amends the standard CGL Coverage Part with both exclusions plus the shared definition — structurally the union of CG 40 48 (Coverage B) and a Coverage A clause.

Critical scope characteristics:

- Triggers only when loss arises from **generative output** (creates content or responses)
- Pre-existing predictive AI, recommendation systems, fraud detection, and other inference-based AI are **arguably outside scope**
- SAMPLE-watermarked PDFs for CG 40 48 and CG 35 08 publicly hosted at assets.alm.com; CG 40 47 verbatim from FC&S Bulletins reproduces operative text identically (see §8)

Style B — Berkley "Absolute" broad exclusion (D&O / E&O / Fiduciary line)

WR Berkley filed **PC 51380 00 (06-24) "Artificial Intelligence Exclusion (Absolute)"** at the management-liability layer. The exclusion is materially broader than ISO's CGL forms — both in what it excludes and in how it defines AI.

Verbatim core exclusion text:

> *"The Insurer shall not be liable to make payment under this Coverage Part for Loss on account of any Claim made against any Insured based upon, arising out of, or attributable to: (1) any actual or alleged use, deployment, or development of Artificial Intelligence by any person or entity, including but not limited to: (a) the generation, creation, or dissemination of any content or communications using Artificial Intelligence; … (f) any alleged representations, warranties, promises, or agreements actually or allegedly made by a chatbot or virtual customer service agent…"*

Verbatim definition of "Artificial Intelligence":

> *"'Artificial Intelligence' means any machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments…"*

This definition captures **non-generative inference**: algorithmic credit-scoring, recommendation systems, fraud-detection AI, decision-support systems. A claim under Berkley's form involving a deterministic ML scoring model with no GenAI component could be excluded.

State filings confirmed: Connecticut. Berkley rolls form approvals state by state; full state map not publicly enumerable.

Other carriers with verified AI exclusion filings

- **AIG (via National Union Fire Insurance Co. of Pittsburgh, PA)** — filings approved in **Idaho and Illinois**; adopted into hospice/home-health program
- **Great American Insurance Group** — filings approved in **Washington**; broader rollout sought

- **Philadelphia Indemnity Insurance Co.** — content-focused exclusion (Coverage B style); covers AI-generated advertising content
- **Hamilton Select Insurance Inc.** — management liability exclusion using "generative artificial intelligence" trigger language

The Big-3 carrier claim — asymmetrically verified

A widely-circulated trade-press claim — that **Berkshire Hathaway, Chubb, and Travelers** are filing AI exclusions in commercial liability lines with **80%+ state regulator approval** — does not hold up uniformly across the three carriers. Verification status as of May 2026:

- **Chubb — moderate confidence.** Independently named by the *Financial Times* in November 2025 alongside AIG, WR Berkley, and Great American — *before* the Wolfe Research analysis ran. So Chubb has TWO independent reporting chains: FT-direct (Nov 2025) and Wolfe/Information (April 2026). FT specifically reports Chubb's approach is to exclude **"widespread/correlated AI loss events"** while continuing to cover discrete AI-related incidents. Specific subsidiary, form number, state, and effective date are not in the public record.
- **Travelers and Berkshire Hathaway — low-moderate confidence, single-sourced.** Both trace only to Wolfe Research's April 2026 review cited in *The Information* (paywalled). Neither is named in FT's Nov 2025 story. Neither appears in any legal-analyst writeup (Hunton, PolicyholderPulse, Lathrop, Wiley) or in Business Insurance / Commercial Risk Online coverage as of May 2026. BHSI's own published material discusses AI as a risk being *monitored*, not coverage being *withdrawn*. We treat the Wolfe/Information claim for these two carriers as reported but unverified.

Practical action for your renewal preparation:

- If your CGL or management-liability carrier is **Chubb**: ask your broker for the AI exclusion endorsement form number and scope (FT corroboration makes it likely something has been filed; specifics undisclosed).
- If your carrier is **Travelers or Berkshire**: ask your broker directly, but the public record doesn't yet confirm specific form filings. Don't assume an exclusion exists until you see the form text.

SERFF Filing Access blocks automated verification at the tool layer. The carriers who appear in independent legal-analyst writeups with named form text (WR Berkley, AIG, Great American, Philadelphia Indemnity, Hamilton Select) are the verified universe; everything else is single-source trade press until your broker confirms.

Coverage line bifurcation — what's at risk where

The most operationally important takeaway:

Insurance line	AI exclusion status	Confidence
CGL Coverage A (BI/PD)	ISO CG 40 47 or 35 08 — narrow, GenAI only	High (verified form text)
CGL Coverage B (Personal & Advertising Injury)	ISO CG 40 48 — narrow, GenAI only	High
CGL Products/Completed Operations	ISO CG 35 08 — narrow, GenAI only	High
D&O / E&O / Fiduciary	Berkley PC 51380 "Absolute" — broad, all inference-based AI	High (verified form text; Berkley confirmed in CT)
Management Liability (private company)	Berkley + Hamilton Select pattern; comparable exclusions emerging	Medium
Cyber	**Moving OPPOSITE direction** — affirmative coverage extensions (Coalition, Cowbell, AXA XL, QBE); draft AI sublimits (Beazley, QBE) NOT YET BOUND on in-force policies. No US cyber AI exclusion form publicly identified.	High (verified negative finding)
Crime / Fidelity	Berkley filed comparable exclusions per beinsure summary	Medium

A customer reading marketing headlines about "AI exclusions" and assuming a uniform exclusion across all insurance lines will be wrong. The risk is line-specific. Your renewal preparation should be line-specific.

Cyber lines specifically — moving in the OPPOSITE direction

The trade-press framing that lumps cyber together with CGL is conflation. As of May 2026, no major US cyber carrier has filed a bound AI exclusion endorsement. What's happening in cyber is structurally different from CGL:

Affirmative AI coverage extensions (the opposite of exclusions):

- **Coalition** — Affirmative AI Endorsement (March 2024), expanded into Active Cyber Policy (April 2025), plus Deepfake Response Endorsement (December 2025)
- **Cowbell** — Prime One (April 2026) with affirmative AI and quantum risk coverage
- **AXA XL** — GenAI Cyber Endorsement covering data poisoning, EU AI Act fines, usage-rights infringement

- **QBE North America** — AI-focused cyber coverages (July 2025) including LLMjacking + regulatory coverage extensions

Draft AI sublimits — NOT YET BOUND on in-force policies:

- **Beazley + QBE** — working on ~10% sublimits (e.g., $250K on $5M cyber tower) per Financial Times reporting (April 2026)
- **Aidan Flynn, Head of Cyber Underwriting Management, Beazley** (FT, April 2026): the sublimit wording is *"still in development and has not yet been applied to in-force policies."*
- Beazley's public position: *no plans to exclude AI*

Underwriting questionnaires asking about AI use — adjusting risk selection (and pricing), not coverage scope.

ISO has NOT released a CY-series AI exclusion to parallel CG 40 47 / 40 48 / 35 08. The CGL AI exclusion infrastructure does not have a cyber-line counterpart in the public record.

The most useful cyber-side framing for your renewal preparation, from a sourced broker quote:

> *"Some insurance carriers come outright to add AI endorsements to clarify they're still intending to cover losses that are initiated by an AI threat actor… Organizations today do not need to panic that their coverage is in a position to deny an AI-related claim if it's for something that is already traditionally intended to be covered by the policy." — Alexandra Bretschneider, VP, Johnson Kendall Johnson (via Insurance Business Magazine, December 2025)*

What this means for your cyber renewal: questions on AI use will appear in your underwriting submission (per Aon, Marsh, Lockton — see §1 quotes), but coverage exclusions specifically targeting AI are not yet in the bound-policy market. The pressure is on **answer quality**, not on **coverage availability** — which is exactly what this brief is structured to help you with.

Definition arbitrage risk — the under-discussed angle

The gap between ISO's narrow definition (machine-based system that **creates content or responses**) and Berkley's broad definition (any machine-based system that **infers... predictions, content, recommendations, or decisions**) is materially load-bearing.

Concrete example: an algorithmic credit-scoring model used by a fintech mid-market customer.

- Under **ISO CG 40 47 (CGL)**: probably outside scope — the model doesn't create generative output
- Under **Berkley PC 51380 (D&O / E&O)**: probably within scope — the model "infers... decisions" from input

This means the same underlying AI system could be **covered under your CGL renewal and excluded under your D&O renewal**. CIOs and General Counsel should compare exclusion definitions across all relevant insurance lines, not just within a single line.

What this means for your renewal preparation

Practical implications for the next 90-180 days before your renewal:

1. **Identify your insurance lines specifically.** CGL, D&O, E&O, Fiduciary, Cyber, Crime — each may have different exclusion exposure.
2. **Ask your broker for the AI exclusion endorsement(s) on file.** Carriers often quietly attach these at renewal without flagging them. Request the form numbers + verbatim text.
3. **Compare definitions across lines.** If your D&O carrier uses Berkley-style "Absolute" language, your exposure is materially different than if all your lines use ISO narrow language.
4. **Inventory your AI footprint by line.** Generative AI (LLM-based content, drafting, code-gen) → most exposed under ISO forms. Predictive ML (scoring, recommendation, classification) → most exposed under Berkley-style forms.
5. **For management liability specifically:** if you're using ML for any decision-influencing function (credit, hiring, claims adjudication, recommendation systems), get the Berkley-style definition into the conversation early. Buying down or negotiating the exclusion may be more economical than discovering it during a claim.
6. **Document your AI governance maturity (10-Layer Assessment).** Even when exclusions are filed, demonstrable governance posture can sometimes get the exclusion modified or buy-down terms offered.

The exclusion landscape is moving fast. This section will be updated as additional filings surface and as SERFF verification work closes the gap on the Big-3-carrier claim.

---

§4 — Example submission answer template

Two-column format. Left: an example question shape from a 2026 mid-market cyber submission. Right: an operator-grade answer template you can adapt with your specific evidence. Bracketed fields `[LIKE THIS]` indicate organization-specific inserts.

Underwriter question	Operator-grade answer template
Do you use generative AI in production? Describe scope + scale.	"Yes. [N] production AI workflows operate across [DOMAIN] use cases. Workflow inventory + ownership matrix maintained in our governance dashboard; available on request under NDA."
What governance framework do you use?	"10-Layer AI Governance framework, cross-mapped to SOC 2 CC6/CC7, EU AI Act risk tiers, NIST AI RMF, ISO/IEC 42001. Framework documentation public at [URL]."
Have you assessed AI risk in last 12 months? Provide documentation.	"Yes. Most recent assessment dated [DATE], scored [X]/100 with documented per-layer remediation roadmap. Assessment document + scoring methodology available on request."
Audit trail for AI decisions — can you produce one?	"Yes. Every action logged with attribution (actor + per-agent identity), input, prompt template version, model version, output, timestamp. HMAC-signed; SIEM-exported. Sample audit trail available on request."
Do AI agents have separate identities from human users?	"Yes. Per-agent identity provisioned via [IDENTITY PROVIDER] with scoped permissions; revocation flows through identity governance platform. RBAC enforced at AI tool layer + integration layer."
Vendor risk register for LLM providers?	"Yes. Register covers [N] AI providers with documented data-handling agreements, residency, SLAs, sub-processor lists. Production workflows model-agnostic; demonstrated multi-provider deployment."
Incident response plan for AI-specific failures?	"Yes. Covers prompt injection, data exfiltration, model jailbreak, audit-trail tampering. Most recent tabletop [DATE] simulated [SCENARIO]. AI incidents tracked separately with named accountable owners + DLQ for failed AI operations."
What's your shadow AI prevention?	"Written AI acceptable-use policy [VERSION], quarterly training cycle. Technical enforcement via SSO restriction of AI tools to managed accounts. Shadow AI detection via [TOOL] with alerting."

These answer shapes are structured to match what brokers report underwriters increasingly expect. They're deliberately concrete (named systems, dated artifacts, available-on-request evidence) rather than hand-wavy. If you can't fill the brackets, that's a gap to remediate before your renewal.

§5 — What we don't claim

In the spirit of operator-honest discipline:

- **We do NOT claim our framework reduces your premium by X%.** No insurer publicly offers vendor-specific governance discounts as of May 2026. Premium impact depends on your carrier, your broker, your full risk profile, and your claims history. Anyone telling you a third-party framework guarantees a premium discount is selling you something.
- **We do NOT claim our framework guarantees coverage adequacy.** AI exclusions are appearing in commercial liability (ISO CGL forms) and reportedly in cyber lines. Documented governance maturity helps you avoid the worst tier; it does not guarantee a specific coverage level.
- **We do NOT claim carrier-side endorsement.** Some carriers will treat our 10-Layer framework as adequate documentation; others may insist on the carrier's own self-assessment tool. Brokers may translate our framework into their preferred format. The framework's value is independent of carrier endorsement.
- **We do NOT recommend dropping any controls you already have.** Adopting the 10-Layer framework is additive — it complements existing SOC 2, ISO 27001, HIPAA, and PCI-DSS programs. It does not replace them.
- **We do NOT promise the questionnaire shape will be stable.** Underwriting questions are evolving quickly. This brief reflects the 2026 question landscape. v2 will follow when meaningful changes emerge.

---

§6 — Three engagement paths

All three are free; the framework is yours regardless of whether you ever work with JieGou.

Path A — Run it yourself (no contact required)

- Download the brief markdown source from [/downloads/cyber-underwriting-readiness-brief-v1.md]
- Use the answer templates in §4 to draft your AI-section response for your next renewal
- Cross-reference the 10-Layer framework at [/10-layer-assessment] for the deeper baseline
- We don't need to hear about it. The framework is operator-honest about not requiring vendor lock-in.

Path B — 45-min walk-through with our team

- Schedule via [/demo] — 30 min discovery, 15 min brief walk-through
- We'll review your specific renewal context (carrier, broker, current submission stage) and identify gaps + quick wins

- No sales pitch. If we're not the right fit, we'll say so and recommend the right shape (broker, SI, or in-house engineering).

Path C — Broker advisory path

- If you'd like us to talk directly with your broker about your AI submission, email partnerships@jiegou.ai with broker name + renewal timing
- We'll join the next packet-prep call as technical advisor (no fee, no commitment)
- Goal: ensure your submission's AI section reflects the operator-grade evidence you actually have

---

§7 — The regulatory landscape (for the audit-defense narrative)

If your CFO or General Counsel asks why this is suddenly relevant:

United States:

- **NYDFS Industry Letter (Oct 16, 2024)** — *Cybersecurity Risks Arising from Artificial Intelligence*. Requires Senior Governing Body oversight of AI cybersecurity programs at NYDFS-regulated entities. Establishes the supervisory expectation that flows to underwriters.
- **NAIC Model Bulletin on AI Use by Insurers** (adopted Dec 2023; 24 states adopted as of 2026) — requires insurers to maintain "Governance with a documented accountability structure... Risk Management and Internal Controls... and Third-Party Vendor Oversight" for AI systems. Insurer-side regulatory floor becomes insured-side underwriting question.
- **NYDFS Industry Letter (Oct 21, 2025)** — third-party service provider risk extension; directly relevant for AI vendor risk register requirements.

United Kingdom + Europe:

- **Lloyd's Market Association** — launched AI Adoption Toolkit (2025) supporting "governance-led implementation across the Lloyd's market." Lloyd's syndicates underwrite a significant portion of mid-market cyber globally.
- **EU AI Act** — high-risk AI systems include underwriting + creditworthiness; creates downstream pressure on EU-presence orgs. Phased enforcement through 2026-27.

Industry standards:

- **NIST AI Risk Management Framework** (US, voluntary) — Govern / Map / Measure / Manage structure increasingly cited in carrier questionnaires.
- **ISO/IEC 42001** AI Management System — emerging international standard; certification path comparable to ISO 27001.

For underwriting purposes, the key point is: **multiple regulators independently require the things your underwriter is asking about.** Your AI governance posture isn't being assessed as a marketing exercise; it's being assessed because regulators are requiring it of your insurer.

---

§8 — Sources

Every claim in this brief is anchored to a publicly verifiable source. Quotes are reproduced from the original publication; URLs are live as of 2026-05-23.

Broker market commentary:

- Aon — *Cyber and E&O: Pricing Holds, but Market Momentum is Shifting* (2026). https://www.aon.com/en/insights/articles/cyber-and-eo-pricing-holds-but-market-momentum-is-shifting
- Aon — *Global 2025 Cyber Risk Report*. https://www.aon.com/cyber-risk-report
- Marsh — *US Cyber Insurance Market Update Q4 2024*. https://www.marsh.com/en/services/cyber-risk/insights/cyber-market-update-q4-2024.html
- Lockton — *Cyber Market Update, December 2025*. https://insights.lockton.com/lockton-market-update/december-2025/cyber
- WTW — *Insurance Marketplace Realities 2026 Cyber Risk*. https://www.wtwco.com/en-us/insights/2025/10/insurance-marketplace-realities-2026-cyber-risk
- WTW — *Emerging AI exposures and the role of cyber and E&O insurance* (Mar 2025). https://www.wtwco.com/en-us/insights/2025/03/emerging-ai-exposures-and-the-role-of-cyber-and-e-and-o-insurance
- Woodruff Sawyer — *2025 Cyber Looking Ahead Guide*. https://www.prnewswire.com/news-releases/woodruff-sawyers-2025-cyber-looking-ahead-guide-highlights-declining-insurance-costs-and-rising-risks-302366421.html

Carrier publications:

- Coalition — *Affirmative AI Endorsement* announcement (Mar 2024). https://www.coalitioninc.com/announcements/coalition-adds-new-affirmative-ai-endorsement-to-cyber-policies
- Coalition — *AI Coverage* page. https://www.coalitioninc.com/ai-coverage
- Cowbell — *Prime One — affirmative AI and quantum coverage* (Apr 2026). https://cowbell.insure/news-events/pr/prime-one-us-emerging-ai-quantum-risks/
- At-Bay — *2025 InsurSec Report*. https://www.at-bay.com/2025-insursec-report/

Regulatory and standards:

- NYDFS — *Industry Letter on Cybersecurity Risks Arising from Artificial Intelligence* (Oct 16, 2024). https://www.dfs.ny.gov/industry-guidance/industry-letters/il20241016-cyber-risks-ai-and-strategies-combat-related-risks
- NAIC — *Model Bulletin on the Use of Artificial Intelligence Systems by Insurers* (April 2024 PDF). https://content.naic.org/sites/default/files/inline-files/AI%20Model%20Bulletin%20-%20April%202024.pdf
- Plante Moran — *How the NAIC AI Model Bulletin Is Evolving* (Mar 2026). https://www.plantemoran.com/explore-our-thinking/insight/2026/03/how-the-naic-ai-model-bulletin-is-evolving
- Lloyd's Market Association — *AI Adoption Toolkit launch*. https://lmalloyds.com/lma-launches-ai-adoption-toolkit-to-support-governance-led-implementation-across-the-lloyds-market/

Trade press:

- Insurance Business — *Cyber insurance enters the AI risk era*. https://www.insurancebusinessmag.com/us/news/cyber/cyber-insurance-enters-the-ai-risk-era-as-limits-wording-and-underwriting-models-shift-565329.aspx
- Reinsurance News — *Lloyd's warns AI is amplifying evolving cyber risk vectors*. https://www.reinsurancene.ws/lloyds-warns-ai-is-amplifying-evolving-cyber-risk-vectors-and-uncertainty-in-coverage-exposure/
- IAPP — *How AI Liability Risks Are Challenging the Insurance Landscape*. https://iapp.org/news/a/how-ai-liability-risks-are-challenging-the-insurance-landscape

AI exclusion filings (§3.5 — added v1.1):

- ISO form CG 40 48 01 26 — *Exclusion: Generative Artificial Intelligence (Coverage B Only)*. Insurance Services Office, © 2025. SAMPLE PDF: https://assets.alm.com/63/68/46ed4bf34a0e807c9695e15c9e19/cg-40-48-01-26-exclusion-generative-artificial-intelligence-coverage-b-only.pdf
- ISO form CG 35 08 01 26 — *Exclusion: Generative Artificial Intelligence (Products / Completed Operations)*. SAMPLE PDF: https://assets.alm.com/3f/6f/918870894682a2e4a733bb0229fd/cg-35-08-01-26-exclusion-generative-artificial-intelligence.pdf
- ISO form CG 40 47 01 26 — *Exclusion: Generative Artificial Intelligence (Coverage A + Coverage B)*. Scope description: Verisk Core Lines emerging-risks article (July 2025); FC&S commentary at PropertyCasualty360 (Oct 2025). https://core.verisk.com/Insights/Emerging-Issues/Articles/2025/July/Week-4/Emerging-Risks-in-ISO-General-Liability-Multistate-Filing
- WR Berkley form PC 51380 00 (06-24) — *Artificial Intelligence Exclusion (Absolute)*. Hunton-hosted verbatim PDF: https://www.hunton.com/assets/htmldocuments/noindex/PC-51380-00-06-24-

Artificial-Intelligence-Exclusion-Absolute.pdf

- Hunton Andrews Kurth — *How Insurance Policies Are Adapting To AI Risk* (Law360 republication; includes Hamilton Select + Philadelphia Indemnity exclusion fragments): https://www.hunton.com/insights/publications/how-insurance-policies-are-adapting-to-ai-risk
- Hunton Insurance Recovery Blog — *The Continued Proliferation of AI Exclusions*: https://www.hunton.com/hunton-insurance-recovery-blog/the-continued-proliferation-of-ai-exclusions
- PropertyCasualty360 / FC&S (Karen Sorrell, Oct 6, 2025) — *General Liability Endorsements: Assault or Battery, Generative AI, Human Trafficking*: https://www.propertycasualty360.com/fcs/2025/10/06/general-liability-endorsements---assault-or-battery-generative-ai-human-trafficking/
- IndependentAgent.com — *Verisk to Roll Out New General Liability Exclusions for Generative AI Exposures*: https://www.independentagent.com/vu_resource/verisk-to-roll-out-new-general-liability-exclusions-for-generative-ai-exposures/
- beinsure.com (summary of *Financial Times* reporting on AIG / Berkley / Great American state-specific filings): https://beinsure.com/news/us-insurers-add-generative-ai-exclusions/
- Trade-press reporting on Berkshire / Chubb / Travelers (Wolfe Research analysis cited in The Information, paywalled; Travelers + Berkshire NOT independently verified beyond this chain; Chubb HAS additional independent corroboration via FT Nov 2025 — see below): https://insuranceintel.substack.com/p/berkshire-chubb-and-travelers-are

Cyber-line specific (v1.2 update — going OPPOSITE direction):

- Insurance Business Magazine — *AI exclusions are creeping into insurance — but cyber policies aren't the issue (yet)* (Dec 2025): https://www.insurancebusinessmag.com/us/news/cyber/ai-exclusions-are-creeping-into-insurance--but-cyber-policies-arent-the-issue-yet-560647.aspx
- Commercial Risk Online — *Beazley has no plans to exclude AI*: https://www.commercialriskonline.com/beazley-has-no-plans-to-exclude-ai/
- AXA XL — *New cyber endorsement extending coverage for Gen AI risks*: https://axaxl.com/press-releases/axa-xl-unveils-new-cyber-insurance-extending-coverage-to-help-businesses-manage-emerging-gen-ai-risks
- QBE North America — *AI-Focused Cyber Insurance Coverages* (July 2025 affirmative coverage, not exclusion): https://www.qbe.com/us/newsroom/press-releases/qbe-north-america-introduces-ai-focused-cyber-insurance-coverages-to-address-emerging-risks
- TechCrunch summary of FT Nov 2025 (independent Chubb corroboration): https://techcrunch.com/2025/11/23/ai-is-too-risky-to-insure-say-people-whose-job-is-insuring-risk/

- Business Insurance — *Insurers, brokers adjust as AI exclusions emerge* (named AIG/Berkley/Great American; cyber side framed as clarifications not exclusions, Bretschneider quote): https://www.businessinsurance.com/insurers-brokers-adjust-as-ai-exclusions-emerge/

CG 40 47 verbatim text (v1.2 update — gap closed):

- PropertyCasualty360 / FC&S Bulletins (Karen Sorrell, Oct 6, 2025) reproduces CG 40 47 operative exclusion clauses verbatim; cross-checked against verified CG 40 48 + CG 35 08 sibling forms for structural + definition consistency. https://www.propertycasualty360.com/fcs/2025/10/06/general-liability-endorsements---assault-or-battery-generative-ai-human-trafficking/

Adjacent framework references:

- /10-layer-assessment — JieGou's published 10-Layer AI Governance Self-Assessment with 30-question scoring
- /reference-architecture — operating substrate detail relevant to audit-trail / per-agent-identity / vendor-portability sections
- /security — JieGou security posture page with cross-framework regulatory mapping
- StackAware — *Cyber Insurance, AI Governance, ISO 42001, and NIST AI RMF*. https://blog.stackaware.com/p/cyber-insurance-ai-governance-iso-42001-nist-rmf — independent operator analysis confirming no formal premium discounts for AI governance as of writing

---

Closing

The cyber insurance market and AI governance are converging at exactly the speed regulators are forcing convergence on insurers. By 2027, AI sections in mid-market cyber submissions will likely be as standardized as MFA questions are today. The orgs that win that transition are the ones that started answering the questions before they were required.

This brief is the document we'd want to have on the desk if we were in your position. The 10-Layer framework + the underwriting categories above are independent of any vendor relationship — they're the operator-grade map of what's coming.

If you want to talk through the renewal-cycle implications for your specific situation, book a 30-min discovery call. If you'd rather run this yourself, the framework + this brief are everything you need.

JieGou — AI Operations Partner for engineering-led mid-market IT

Cyber Underwriting Readiness Brief · v1.2 · 2026-05-23 · Public artifact · operator-grade not marketing-grade · no sales-call condition · sources at §8

v1.1 expansion (2026-05-23): added §3.5 "AI Exclusion Landscape" with verbatim ISO and Berkley form text, two-style framework (narrow CGL vs broad management-liability), coverage-line bifurcation table, and definition arbitrage analysis.

v1.2 update (2026-05-23, same-day): verification queue closeout. Added CG 40 47 verbatim text from FC&S Bulletins; reshaped Big-3 carrier framing to asymmetric verification (Chubb has FT independent corroboration; Travelers + Berkshire remain single-sourced to Wolfe/Information); added cyber-line "moving opposite direction" subsection with Beazley + QBE draft-sublimit context + Bretschneider quote framing cyber as no-exclusions-yet category. Sources updated.