

Insurable Digital Labor

The operating infrastructure for underwriting AI-mediated work

Fabio A. Thiers, MD, PhD

ConverGen AI, New York, NY, USA

ABSTRACT

Business insurance is established operating infrastructure for healthcare. AI, however, became an operational risk faster than it became a describable insurance exposure. Through 2025 and 2026, the market moved in two directions on this matter. Standardized AI exclusions entered commercial general liability, while a specialist market formed around AI products, vendors, and certified agents. Neither movement fully addresses the exposure that matters most to a deploying organization. Once a health system supplies the role, context, credentials, workflow, and authority, the risk belongs to the operation rather than any product. This article argues that the tractable underwriting unit is the operation of Digital Labor in context, that agentmatics can supply underwriting evidence before credible loss history exists, and that the practical market pathway runs through endorsements on existing institutional programs. Insurability is best read as an external test of whether AI-mediated work has come under management.

Keywords: healthcare AI; Digital Workers; Digital Labor; Digital Operations; DOpS; silent AI; agentmatics; operational telemetry; risk financing; insurance underwriting

Key points

- **AI became operational risk before becoming insurable.** Silence is now ending through exclusion at least as fast as through affirmative coverage of AI-related risk.
- **The tractable underwriting unit is the operation of Digital Labor in context.** It is the bounded Digital Worker acting under organizational authority, not the foundation model and not every locally configured AI artifact.
- **Product-level cover and agent certification solve a different problem.** Both center underwriting on a specific model, product, or agent, and neither scales readily to a proliferating portfolio of locally configured roles.
- **Digital Operations supplies the persistent operating record.** Underwriting, authority decisions, and claims reconstruction all draw on that record, which can support Digital Labor insurability.

Business insurance finances operational risk

Healthcare organizations do not wait until operations are completely risk-free to take action. They control what they can and cover the remaining risks through various insurance types, including professional and hospital liability, cyber and privacy policies, errors and omissions, crime coverage, excess layers, captives, and self-insured retentions. Insurance does not eliminate operational risk but helps make the financial consequences of remaining risks predictable, enabling institutions to operate, borrow, contract, and expand.

The arrangement works because insurers can characterize an organization and what it does. They do not write a separate policy for every employee, workflow, or piece of software. They underwrite defined operations, exposure bases, controls and loss experience, and a small set of familiar denominators does the translation. Payroll, beds, encounters, and revenue turn an institution into something an underwriter can price.

Liability analysis for healthcare AI was already difficult when these systems only advised clinicians, and the idea that liability insurance might itself drive responsible adoption predates the current wave.^{1,2} What is different now has little to do with model opacity. AI is becoming embedded in real work faster than organizations can inventory, bound, or describe it. The pressing question is not whether an AI product is insurable. It is whether an insured organization can explain how AI participates in its

operations, what authority it holds, how much work it performs, which controls apply, and what happened when something failed.

AI became a risk before becoming insurable

Healthcare AI has moved decisively from advice toward work, and deployment has outrun governance. A 2026 report from the Center for Connected Medicine and KLAS Research found that more than 90 percent of surveyed health systems had deployed third-party AI, while only 44 percent maintained dedicated testing environments and 63 percent called their AI strategy developing or *ad hoc*.³ Individual use has risen just as fast, with 81 percent of physicians reporting professional AI use in 2026 against 38 percent in 2023.⁴ Unsanctioned use compounds the picture.⁵

The resulting exposure is fragmented in a particular way. A prompt gets reused across teams without a registered owner. A browser agent operates through a shared service account. A model update changes behavior with no role-level record. Logs preserve API calls but not institutional purpose, supervisor, approval, or outcome. The organization ends up with useful AI-mediated work and no coherent operating record of it.

Insurance markets are meeting the same ambiguity from the other side. Brokers use the term silent AI for AI-related loss that falls under policies never drafted or priced for it.^{6,7} Verisk's Insurance Services Office published three generative-AI endorsements carrying January 2026 edition dates, addressing bodily injury and property damage, personal and advertising injury, and products and completed operations.⁸ Individual carriers have filed proprietary language too, including an absolute AI exclusion for directors and officers, errors and omissions, and fiduciary lines.⁹ Coverage counsel describe the result not as clarity but as fragmentation across lines, with gap risk where no single policy responds.¹⁰ Silence is ending through subtraction at least as fast as through affirmative coverage of AI-related risk.

Healthcare-specific reviews consistently find that most broker analysis reports indicate a lack of mention of AI in healthcare errors and omissions language. Some carriers are starting to exclude unregulated AI applications in documentation and clinical decision support, and early AI liability endorsements are emerging. These endorsements often apply only when a tool reaches a specific validation or certification standard.¹¹ That detail matters, as the market is saying, in its own idiom, that coverage will follow evidence.

Accountability moving to deploying organization

Two developments in the first half of 2026 point the same way. Neither comes from the insurance market itself.

The first is accreditation. On 1 June 2026, the Joint Commission launched its Responsible Use of AI in Healthcare certification, organized around governance, data management, risk and bias reduction, monitoring and validation of safety performance, and transparency and training. It recognizes organizations and explicitly does not validate or certify individual AI products.¹² Its unit of assessment is therefore the organization, not the individual AI product.

Reinsurance analysis shows that Gallagher Re, in collaboration with MIT and Testudo, found AI-native liabilities do not fit well in traditional categories. They also noted that courts and regulators tend to hold deployers responsible rather than vendors, and that liability caps for vendors leave deployers vulnerable. The same analysis warns that concentration among a few foundation-model providers could propagate one defect across many unrelated policyholders.¹³ Together, these point to something the market has not yet operationalized. Responsibility is landing on the deploying organization, and assessment is moving there too, but no one has specified what the organization must actually run to be assessable. This article addresses that gap.

The problem is in the definition of unit of risk

Product-level AI insurance deserves credit in this setting, not dismissal. It solves a different problem and solves it well. Munich Re has underwritten AI model performance since 2018 through aiSure, settling claims on measurable performance data rather than conventional loss adjustment, and extended distribution through Mosaic in 2026.^{14,15} Armilla, a Lloyd's coverholder, writes

affirmative AI liability for underperformance, hallucination and drift, and with Chaucer pairs dedicated AI limits with cyber limits under predefined allocation rules.¹⁶ The Artificial Intelligence Underwriting Company's AIUC-1 standard puts agents through adversarial evaluation with independent audit and quarterly retesting, and the first policy written against it was bound in February 2026.^{17,18} These are real advances, and together they establish that structured evidence can unlock coverage before loss experience exists.

They also share one structural feature. Each is centered on underwriting an artifact, whether a model, a product, or an agent. Munich Re has extended performance cover to models an organization develops for itself, which is the important counterexample, but it proceeds on the same basis: technical due diligence of one specific model against defined performance thresholds.¹⁹ AIUC-1 certifies a particular agent and is retested quarterly because agent behavior changes quickly. Both approaches scale with the number of artifacts assessed.

Now consider what health systems are beginning to build. They can assemble a prior-authorization Digital Worker from a foundation model, payer-policy retrieval, a portal harness, an EHR service account, and a named supervisor. Then they can add a referral worker, a denial-appeal worker, and a scheduling worker. Each is locally configured and changes when a payer policy or a model version changes. No single vendor owns or can certify the complete local operating configuration because the organization assembles much of it. It is implausible that every such role will receive its own policy and its own technical due diligence, or that insurers will build capacity to evaluate thousands of transient local configurations one at a time.




FIGURE 1 - Where the insured exposure actually lives, and how accountability divides across the organization, the management platform and the insurer.

The regulatory boundary does not resolve the situation. Software intended to support a healthcare facility's administration is generally not a medical device, and the FDA's revised clinical decision support guidance hinges on whether a licensed professional can independently review the basis for a recommendation.^{20,21} Therefore, an AI system can cause financial, privacy, or patient-access losses without triggering device regulation.

Quanyan Zhu's 2026 analysis of agentic-AI insurance offers the most complete answer so far on the carrier side. It proposes a layered ecosystem of complementary coverages with explicit allocation rules, in which policies respond only to scheduled AI systems whose approved tools and permissions have been declared.²² That architecture is sound, and this article assumes it rather than disputing it. It does presuppose something on the insured's side that does not yet exist as a discipline. Someone has to maintain the schedule, keep it current as roles change, and produce the evidence when a loss occurs.

The durable formulation is simple enough. Insure products for product risk and underwrite organizations for operational risk. Once an organization supplies the role, the context, the credentials, the workflow, the supervisor, and the authority, the exposure belongs to the operation. That operation is a stack, and every layer can reduce risk or create a failure path

From scattered AI to managed Digital Labor

The management problem becomes tractable once the unit shifts from technology to work. Earlier articles developed the vocabulary, and the insurance argument needs only its outline.

A Digital Worker is a bounded, role-bearing unit through which AI-enabled capability performs defined work under organizational authority. The designation implies no personhood and no employment status. It creates a persistent management object to which purpose, owner, supervisor, identity, permissions, context, evidence obligations, performance, incidents, authority, and lifecycle can all be attached. Digital Workers collectively supply Digital Labor. Digital Operations is the management discipline for designing, supervising, measuring, and improving that labor, and a Digital Operations System (DOPs) is the infrastructure that applies it consistently.

The insurance value of this framing is continuity. Models, prompts, tools, vendors and interfaces change constantly. The institutional role does not have to change with them, and the record of what work is authorized, who supervises it, and what has happened need not vanish when a component is swapped out.

Digital Operations creates a stable management structure around unstable technology, which is what an underwriter needs. The carrier does not have to certify that a model will remain safe for every purpose. It has to understand what the insured organization is doing as AI capabilities change.

What the Underwriter Can Actually Read

The same AI capability described two ways. Only one of them can be priced.

Scattered AI



NO INVENTORY. NO DENOMINATOR.

Nothing to describe, bound or investigate.

Managed Digital Labor

Registry Named roles, owners, authority

Exposure Measured consequential actions

Controls Identity, least privilege, change

Behavior Overrides, escalations, incidents

OBSERVABLE. BOUNDED. RECONSTRUCTABLE.

Visible enough to supervise, bounded enough to price.



The unit of underwriting is the work being performed, not the software performing it.

FIGURE 2 - The same AI capability described two ways. Only one of them gives a carrier something to price.

Waiting for actuarial history is not a strategy

Historical loss experience remains fundamental to insurance, and nothing here suggests otherwise. The difficulty is that Digital Labor will not become a stable actuarial object soon enough for loss history to lead. Models improve on a cadence measured in months, and prompts, policies, permissions, and authority change continuously. One model supports thousands of heterogeneous local implementations, and a shared model or configuration can produce correlated failures across many acts at once. Furthermore, exposure is non-stationary, thin in credible data, and structurally dependent, which limits the usefulness of experience rating alone.

None of this means actuarial methods fail. It means the method needs inputs other than history while history accumulates. Cyber insurance is the closest mature analog, having developed through exposure assessment, control review, scenario analysis, and accumulation management well before credible loss triangles existed.²³ Zhu reaches the same conclusion from first principles, arguing that insurers cannot wait for those triangles because market expectations and adverse-selection dynamics will be settled long before the data matures.²² The practical substitute for missing history is not certainty. It is contemporaneous operational evidence: how much digital work is occurring, whether safeguards were operating, what the workers did, and how quickly problems were detected and corrected.

Agentmatics is the evidence layer

Agentmatics is operational telemetry for Digital Labor. It is the versioned, tamper-evident record of what Digital Workers were authorized to do, what they did, and under which controls. The name matters less than the function, and the function follows from three questions that any underwriter or claims investigator will eventually ask.

How much Digital Labor is being performed? Exposure metrics translate Digital Labor into rating bases insurers already understand. These include worker-hours by job family, privileged action units that weight consequential acts such as record writes,

claim submissions, and payment instructions, and communication and data-access units. These are the payroll analogs for a digital workforce.

Were safeguards operating? Control signals record machine identity and least privilege, step-up authorization and dual control, change control and version pinning, monitoring status, review sampling, and audit-log integrity. ISO/IEC 42001 and the NIST AI Risk Management and Cybersecurity Frameworks give useful structure here.²⁴⁻²⁶ But certifications and annual attestations are static, and underwriting a moving operation requires proof that controls operated between attestations.

What did the workers actually do? Behavioral signals capture overrides, escalations, near misses, drift indicators, tool errors, and recovery actions. Most consequentially, they capture escaped material errors, meaning those not prevented, detected, corrected, or escalated before they produced consequence.

Taken together, these signals make a claim file reconstructable. Telemetry as an underwriting input is an established pattern, not a speculative one. Usage-based auto insurance rates observed driving, cyber underwriting draws on continuous security posture, and pay-as-you-go workers' compensation reconciles premium to actual payroll.^{27,28}

Digital Labor carries an advantage none of these had, because the relevant acts and controls can usually be instrumented at greater granularity than human work. Proposed agentic-AI policy architectures already contemplate preservation of prompts, tool traces, access logs, version records, and approvals as conditions of coverage.²² This points toward evidence becoming contractual rather than merely managerial.

An embedded always-on Dynamic Insurability Score can summarize this record, provided it stays explainable, versioned, reproducible, and subordinate to the underlying evidence. It is a signal. It is never a guarantee of safety or a proof of coverage.




FIGURE 3 - Three underwriting questions, the signal families that answer them, and the record they feed.

Keep authority inside the underwritten envelope

The most useful role for evidence is operational rather than contractual. Neither insurers nor insureds want coverage mechanics that fluctuate hour by hour. The design proposed here keeps policy terms stable through the policy period while the management system narrows, pauses, or expands higher-risk authority, so that the operation stays inside the conditions both parties accepted. The bands described below are an operating architecture rather than an existing market convention, and the principle matters more than the taxonomy.

Property insurance offers a cautionary precedent. Protective-safeguards endorsements have long required sprinklers or alarms to be operational as a condition of coverage, and the logic maps onto logging, drift monitoring, and dual control. Strict warranties, though, create coverage-denial traps that would poison an early market. The safer posture treats controls as eligibility and authority gates, reserving denial for clear misrepresentation or willful disabling of safeguards.




FIGURE 4 - Four authority bands, what each permits, and the coverage posture and evidence that go with it.

Endorse the existing insurance tower first

The practical way into this market is coverage attached to an institution's existing program rather than stand-alone policies for individual Digital Workers. That preserves familiar underwriting relationships. The carrier continues to underwrite the organization but now receives a far clearer description of the Digital Labor operating inside it.

Coverage is always policy-specific, so what follows is a mapping, not an interpretation. Technology errors and omissions is the natural line through which to address administrative digital work such as scheduling, registration, and authorization preparation, once a worker registry, telemetry, and validated scope limits exist. Cyber and privacy may respond to PHI handling, AI security events, and agent-mediated disclosure, conditioned on identity architecture, least privilege, and audit-grade logging. Crime and social engineering may respond to fraudulent instructions and payment diversion, conditioned on dual control over funds. Professional liability would reach clinical support and patient-facing actions, and should wait for proven claims forensics, scope limits, event aggregation and staged authority. A coordination endorsement could define how these respond together whenever two or more of them touch Digital Labor.

Administrative Digital Labor should lead, because it can be bounded, measured, supervised, and reconstructed without direct clinical decision-making. Clinical work should follow more cautiously, once the discipline has been demonstrated in production.

Coordination matters as much as coverage, because one incident can cross policy boundaries and be pleaded as professional negligence, technology error, privacy breach or fraud. Specialist markets already have coordinated structures.¹⁶

A plausible pricing progression follows the evidence. At first, evidence would inform eligibility, exclusions, sublimits, retentions and renewal negotiation, without a dynamic premium. Next, annual policies might carry a deposit premium with an audit true-up on measured exposure, as payroll audits work in workers' compensation. Periodic adjustment within pre-agreed bands should come later, and first in captives and excess-and-surplus structures. It would not be advisable to introduce telemetry-style pricing before telemetry-grade evidence exists, as it would reproduce silent AI under a subtler name.

Digital Labor Events and correlated failure

The exposure that most distinguishes Digital Labor from human labor is correlation. A stale payer policy, a prompt change, a model update, a permission defect, or a workflow rule can generate hundreds of similar acts before anyone notices. Machine speed turns configuration error into accumulation risk.

Concentration can compound the risk. Reinsurance analysis indicates that dependence on a limited number of foundation-model providers could spread a single flaw across many unrelated policyholders. In healthcare, a similar pattern occurs, with a national analysis revealing that three ambient documentation products are used by over 80 percent of adopting hospitals. While shared dependency doesn't automatically prove correlated insured loss, it creates the scenario where such loss can happen. Portfolios assuming independence may not anticipate this risk.

Policy and operating design should therefore recognize a Digital Labor Event, meaning related acts aggregated around a common root cause, version or time window for notice, retention, limits, remediation and reinsurance. The cyber analogy is instructive. One breach affects many records and is still investigated as a single event file, with standardized evidence requirements and pre-agreed forensic workflows.

Aggregation also disciplines operations, because the insurer's accumulation risk is bounded by the insured's blast-radius controls. Eligibility conditions should require rate limits, caps on high-volume action types, staged rollouts, version pinning, rollback, and rapid credential revocation. Sublimits for high-correlation perils such as mass messaging can relax as change control and telemetry prove stability. Correlation is not an argument against Digital Labor. It argues for engineering the operation, so one defect stays a bounded event.

What a Digital Operations System must provide

None of this necessitates a healthcare organization to develop its own models, infrastructure, or integrations. Instead, the organization can obtain foundation models, cloud services, connectors, and external expertise. What must remain under the organization's control are the elements that define accountability. If role logic, context, corrections, authority decisions, and operational history are stored within opaque vendor systems, the organization retains liability but loses the ability to explain how its work was carried out. This situation is not vendor lock-in but institutional hollowing. An effective management system, whether called a Digital Operations System or otherwise, should provide at least the following features.

- **Define.** A registry of roles rather than tools, holding every Digital Worker with its purpose, owner, supervisor, job family, permitted and prohibited actions, authority band and lifecycle state. A persistent Digital Worker Record that carries role definition, context sources, validation evidence, performance, incidents, and authority changes as one institutional object. A distinct non-human identity per worker, with permissions scoped to the role. A managed context layer of authoritative local knowledge, with explicit freshness and ownership.

- **Control.** Execution boundaries and blast-radius limits: confirmations, action gates, rate limits, staged rollout, version pinning and rollback. Supervision that is substantive rather than ceremonial, with defined review, escalation, approval and override. Authority management tied to evidence, giving a defensible basis for expanding, narrowing, pausing or revoking what each worker may do.

- **Evidence.** Telemetry by default, capturing inputs, sources, tool calls, decisions, approvals, actions, outcomes and corrections in versioned, tamper-evident form. Incident review and change control that classify failures, near misses and escaped material errors, with rules for when a change requires revalidation or a narrowing of authority.

- **Preserve.** Portability and exportability, so the organization can produce an insurer-readable and auditor-readable account of its workforce without reconstructing it from several vendors after a loss.

Run in sequence, these capabilities describe a working cycle rather than a project. A role is registered, validated locally before authority is granted, then constrained through identity and blast-radius limits. It operates with agentmatics on every consequential act, holds authority inside the underwritten envelope, attracts coverage by endorsement, and contributes event and renewal data that re-enters validation on the next pass. A system meeting that specification answers the underwriter's questions as a by-product of running the operation. A system missing several will not, no matter how capable its models are.

The solution has to scale down as well as up

The obvious objection is that this describes infrastructure only a large health system could maintain. It deserves a serious answer, because a management system that requires sophistication as a prerequisite has failed at its own purpose. Such a system should reduce the expertise required to operate safely, not raise it. A small specialty practice may run three Digital Workers, one operational supervisor, and a handful of workflows, and it still benefits from the same structure at a smaller scale. The Joint Commission's certification designers made a similar judgment, describing standards as adaptable rather than prescriptive, so a small rural organization could meet them without a dedicated governance department.^{12,30}

Embedded AI advisory is what makes this practical. The management system can meet an organization where it actually is. This type of workflow-embedded solution can interview operators about how work is really performed, translate that into role and context requirements, identify missing controls, guide shadow validation, maintain the Digital Worker Record, and prepare evidence an underwriter can read.

One reinforcing dynamic is worth naming plainly. The same capability that makes locally configured Digital Labor easy to create also makes the discipline needed to control it far cheaper to operate. For many recurring operational roles, running Digital Labor through one institutional system will become simpler than maintaining an expanding collection of disconnected deployments of AI solutions.

Total cost of risk, and what insurance does not do

The economic case should not depend on lower premiums. Much healthcare liability risk is retained through deductibles, self-insured retentions and captives, and more is being retained as capacity contracts and severity rises.^{31,32,33}

The right lens is total cost of risk, which includes retained losses, defense costs, downtime, remediation, and premium. Earlier detection of correlated failure, evidence that shortens disputes, and verifiable renewal data in place of vague questionnaires can all create value before premium changes.

Insurance finances residual risk under defined terms. It does not replace the regulatory and governance obligations that govern the

underlying work. The HIPAA Security Rule requires safeguards for electronic protected health information, and Section 1557 requires covered entities to identify and mitigate discrimination risk from patient-care decision-support tools.^{34,35} State law is moving quickly, with 2026 statutes constraining autonomous clinical decision-making and limiting reliance on AI in utilization review.³⁶ Digital Workers hold no licenses, so boards regulate the humans and institutions using them.³⁷ Underwriters should read this stack as a floor: insurance sits above governance, never in place of it.

The executive test

Before seeking affirmative coverage, an organization should be able to answer the questions an underwriter or a claims investigator will eventually ask. None requires a new theory. Each asks only whether the operation keeps a record of itself, and a managed operation answers in minutes rather than meetings.



FIGURE 5 - The executive test. Seven questions an underwriter or a claims investigator will eventually ask, and the evidence each one expects.

Close: the window is narrower than it looks

Insurance markets do not require zero risk. They require risk that can be described, bounded, priced, monitored and investigated. Digital Labor meets that standard when Digital Operations produces claims-grade evidence of scope, controls, supervision, telemetry, incidents and recovery.

There is urgency to develop and broadly deploy solutions to manage risk emanating from Digital Labor. Exclusions are already attaching at renewal, to policies written before most organizations could describe their AI use at all. The key issue is that adoption is not pausing while the market settles. The specialist market now forming is organized around artifacts and their suppliers, because artifacts are what can currently be assessed. If that becomes the settled architecture, deploying organizations will hold a growing pool of operational exposure that no product-level instrument was designed to reach.

In the meantime, the risk is a widening mismatch. AI-mediated work expands while insurers retreat from exposures they cannot characterize, leaving organizations with more retained or ambiguous risk precisely as their dependence on AI grows. The alternative is that the deploying organization becomes describable, with its roles named, its authority bounded and its incidents reconstructable, so that coverage can follow the evidence as it did in cyber. What separates the two is not a model capability. It is whether organizations build the management layer while they still hold the initiative.

None of this removes risk, and no score, endorsement, or telemetry stream should ever be described as if it did. The pathway offers something narrower and more valuable. It offers residual risk made visible enough to supervise, bounded enough to price, and documented enough to defend. An organization that reaches that state has not merely become insurable. It has demonstrated that its Digital Labor is under management, which is the condition on which everything else, including the insurance, depends.

Glossary of key concepts

- **Silent AI exposure.** AI-related losses that may fall under existing policy wordings never drafted or priced for AI, creating uncertainty about coverage, exclusions and claims handling.
- **Digital Worker.** A bounded, role-bearing unit through which AI-enabled capability performs defined work under organizational authority. Not a legal person and not an employee.
- **Digital Operations System (DOPs).** The management infrastructure through which Digital Workers are created, contextualized, credentialed, supervised, measured, improved and retired.
- **Agentmatics.** Operational telemetry for Digital Labor, comprising exposure metrics, control signals and behavioral signals captured in versioned, tamper-evident form to support management, underwriting, claims review and improvement.
- **Digital Labor Event.** Related Digital Worker acts aggregated as one event for retentions, limits, notice, remediation, and reinsurance, usually bounded by a common root cause, version, or time window.
- **Escaped material error.** An error not prevented, detected, corrected or escalated before it produced consequence.
- **Total cost of risk.** Retained losses, defense and remediation costs, downtime, and premium considered together, which is the appropriate lens for evaluating insurable Digital Labor.

Live Article information

This PDF is a dated snapshot of a continuously maintained Live Article. Material updates undergo editorial review, and substantive public revisions are released when warranted.

The current maintained article and evidence base are available through the ConverGen Digital Operations Advisor: www.convergen.ai/advisor/

Peer-review status: This public edition has not undergone external peer review.

Competing interests: The author is the founder and CEO of ConverGen AI, which develops and commercializes a Digital Operations System (DOPs) for healthcare organizations.

AI-assisted tools: AI-assisted tools supported research synthesis, drafting, editing, and figure production. The author reviewed and takes responsibility for all content.

Author ORCID: 0009-0008-1814-1983.

References

- Mello, M. M. & Guha, N. Understanding liability risk from using health care artificial intelligence tools. *N. Engl. J. Med.* 390, 271-278 (2024). <https://doi.org/10.1056/NEJMhle2308901>
- Stern, A. D., Goldfarb, A., Minssen, T. & Price, W. N. II. AI insurance: how liability insurance can drive the responsible adoption of artificial intelligence in health care. *NEJM Catal. Innov. Care Deliv.* 3 (2022). <https://doi.org/10.1056/CAT.21.0242>
- Center for Connected Medicine & KLAS Research. Health system AI adoption and governance (2026). <https://connectedmed.com/resources/>. Accessed 19 Aug 2026.
- American Medical Association, Center for Digital Health and AI. 2026 Physician Survey on Augmented Intelligence (March 2026). <https://www.ama-assn.org/system/files/physician-ai-sentiment-report.pdf>. Accessed 19 Aug 2026.
- Wolters Kluwer Health. Shadow AI in healthcare (2026). <https://www.wolterskluwer.com/en/expert-insights>. Accessed 19 Aug 2026.
- Lior, A. & Madhok, S. Insuring the AI age. *WTW* (2025). <https://www.wtwco.com/en-us/insights/2025/12/insuring-the-ai-age>. Accessed 19 Aug 2026.
- The Geneva Association. Gen AI Risks for Businesses: Exploring the Role for Insurance (2025). <https://www.genevaassociation.org/publication/digital-ai-transformation/gen-ai-risks-businesses-exploring-role-insurance>. Accessed 19 Aug 2026.
- Insurance Services Office (Verisk). Generative artificial intelligence exclusions for commercial general liability: CG 40 47 01 26, CG 40 48 01 26 and CG 35 08 01 26 (2026). Filed forms; available to subscribing insurers.
- Jergler, D. Insurer interest in AI exclusions growing as risk becomes omnipresent. *Insurance Journal* (17 Aug 2026). <https://www.insurancejournal.com/magazines/mag-features/2026/08/17/881424.htm>. Accessed 19 Aug 2026.
- Fenwick & West LLP. The end of silent AI? Emerging AI exclusions, coverage fragmentation, and practical implications for policyholders (2026). <https://www.fenwick.com/insights/publications/end-silent-ai-emerging-ai-exclusions-coverage-fragmentation-and-practical-implications>. Accessed 19 Aug 2026.
- EPIC Insurance Brokers & Consultants. DoNoHarm.exe: the liability reckoning for AI in U.S. healthcare (2026). <https://www.epicbrokers.com/insights/donoharm-exe/>. Accessed 19 Aug 2026.
- The Joint Commission. Responsible Use of AI in Healthcare (RUAIH) Certification (2026). <https://www.jointcommission.org/en-us/certification/responsible-use-of-ai-in-healthcare>. Accessed 19 Aug 2026.
- Gallagher Re, Massachusetts Institute of Technology & Testudo. Smart Systems, Blind Spots: Rethinking Insurance for the AI Era (2026). <https://www.ajg.com/gallagher/news-and-insights/smart-systems-blind-spots-rethinking-insurance-for-the-ai-era/>. Accessed 19 Aug 2026.
- Munich Re. aiSure: insuring the performance of artificial intelligence solutions (2026). <https://www.munichre.com/en/solutions/for-industry-clients/insure-ai.html>. Accessed 19 Aug 2026.
- Mosaic Insurance. Mosaic partners with Munich Re aiSure to provide coverage for AI vendors (2026). <https://www.mosaicinsurance.com/resources/press-releases/>. Accessed 19 Aug 2026.
- Chaucer Group & Armilla AI. Vanguard AI: a coordinated AI liability and cyber insurance structure (2026). <https://www.chaucergroup.com/news/>. Accessed 19 Aug 2026.
- Artificial Intelligence Underwriting Company. AIUC-1: a certification standard for AI agents (2026). <https://www.aiuc-1.com/>. Accessed 19 Aug 2026.
- ElevenLabs. ElevenLabs secures first-of-its-kind AI agent insurance (Feb 2026). <https://elevenlabs.io/blog/aiuc-announcement>. Accessed 19 Aug 2026.
- Munich Re. Insure the performance of self-developed artificial intelligence solutions (2026). <https://www.munichre.com/en/solutions/for-industry-clients/insure-ai/ai-self.html>. Accessed 19 Aug 2026.
- U.S. Food and Drug Administration. Is the software function intended for administrative support of a health care facility? (2022). <https://www.fda.gov/medical-devices/digital-health-center-excellence/step-2-software-function-intended-administrative-support-health-care-facility>. Accessed 19 Aug 2026.
- U.S. Food and Drug Administration. Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff (2026 revision). <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>. Accessed 19 Aug 2026.
- Zhu, Q. Insurance of agentic AI. Preprint at arXiv (2026). <https://doi.org/10.48550/arXiv.2606.05449>
- National Association of Insurance Commissioners. Cybersecurity Insurance Report (2025). https://content.naic.org/sites/default/files/inline-files/2025_Cybersecurity_Insurance%20Report.pdf. Accessed 19 Aug 2026.
- International Organization for Standardization. ISO/IEC 42001:2023 Information Technology, Artificial Intelligence, Management System (2023). <https://www.iso.org/standard/81230.html>. Accessed 19 Aug 2026.
- National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1 (2023). <https://doi.org/10.6028/NIST.AI.100-1>

26. National Institute of Standards and Technology. The NIST Cybersecurity Framework (CSF) 2.0. NIST CSWP 29 (2024). <https://doi.org/10.6028/NIST.CSWP.29>
27. National Association of Insurance Commissioners. Consumer insight: understanding usage-based insurance (2023). <https://content.naic.org/article/consumer-insight-understanding-usage-based-insurance>. Accessed 19 Aug 2026.
28. Bitsight & Gallagher Re. Scanning the Horizon: cybersecurity analytics and cyber claims (2024). <https://www.bitsight.com/resources/gallagher-res-scanning-horizon>. Accessed 19 Aug 2026.
29. Yang, F. et al. Ambient artificial intelligence tool adoption in US hospitals and associated factors. *Am. J. Manag. Care* 32, e25-e30 (2026). <https://doi.org/10.37765/ajmc.2026.89876>
30. Grubbs, K. & Walders, W. Designing an adaptable AI certification for organizations of every size. *Healthcare Dive* (18 Jun 2026). <https://www.healthcaredive.com/news/joint-commission-adaptable-ai-blueprint-certification-ken-grubbs-william-walders/823201/>. Accessed 19 Aug 2026.
31. WTW. Insurance Marketplace Realities 2026 Spring Update: Healthcare Professional Liability (2026). <https://www.wtwco.com/en-us/insights/2026/05/insurance-marketplace-realities-2026-spring-update-healthcare-professional-liability>. Accessed 19 Aug 2026.
32. Milliman. 2025 Medical Professional Liability Update (2025). <https://www.milliman.com/en/insight/2025-medical-professional-liability-update>. Accessed 19 Aug 2026.
33. HCA Healthcare. Annual Report on Form 10-K for the Year Ended 31 December 2024. U.S. Securities and Exchange Commission (2025). <https://www.sec.gov/Archives/edgar/data/860730/000095017025020134/hca-20241231.htm>. Accessed 19 Aug 2026.
34. U.S. Department of Health and Human Services. Summary of the HIPAA Security Rule (2024). <https://www.hhs.gov/hipaa/for-professionals/security/laws-regulations/index.html>. Accessed 19 Aug 2026.
35. U.S. Department of Health and Human Services. Nondiscrimination in Health Programs and Activities: Final Rule. *Fed. Regist.* 89, 37522-37703 (2024). <https://www.federalregister.gov/documents/2024/05/06/2024-08711/nondiscrimination-in-health-programs-and-activities>. Accessed 19 Aug 2026.
36. Holland & Knight LLP. States continue efforts to regulate AI in healthcare: a review of legislation passed in 2026 (2026). <https://www.hklaw.com/en/insights/publications/2026/05/states-continue-efforts-to-regulate-ai-in-healthcare>. Accessed 19 Aug 2026.
37. Federation of State Medical Boards. Navigating the Responsible and Ethical Incorporation of Artificial Intelligence into Clinical Practice (2024). <https://www.fsmb.org/siteassets/advocacy/policies/incorporation-of-ai-into-practice.pdf>. Accessed 19 Aug 2026.