Addendum to The CAS AI Primer:

Practical Guidance for Actuaries

By the CAS Artificial Intelligence Working Group

Casualty Actuarial Society

October 2026

About the Casualty Actuarial Society and the Artificial Intelligence Working Group

The Casualty Actuarial Society (CAS) is a leading international organization for credentialing and professional education. Founded in 1914, the CAS is the world's only actuarial organization focused exclusively on property and casualty risks and serves over 11,000 members worldwide. CAS members are experts in property and casualty insurance, reinsurance, finance, risk management, and enterprise risk management. Professionals educated by the CAS empower business and government to make well-informed strategic, financial and operational decisions.

The CAS Artificial Intelligence Working Group was established to fulfill the CAS mission to advance the body of knowledge on a technology that is transforming actuarial practice. Its objective is to encourage the exploration of artificial intelligence in actuarial practice through research to help educate members, build knowledge, provide practical insight, and establish CAS members as thought leaders.

Recommended Citation:

CAS. 2026. *Addendum to The CAS AI Primer: Practical Guidance for Actuaries*. 2nd ed. Arlington, VA: The Casualty Actuarial Society.

© 2026. Casualty Actuarial Society.

Acknowledgements

The CAS would like to thank Shine Wang for her ongoing leadership to develop the first and subsequent editions of The CAS AI Primer. Contributors and reviewers of this addendum include Emma Chen, Mario DiCaro, Nick Easley, Brooke Engle, Nancy Hu, Bobby Jaegers, Lily Liang, Vajira Manathunga, Arnav Vashishth, and Shine Wang.

Contents

Introduction 1

1. US AI Regulation and Its Impact on Actuarial Practice 2

2. Excel CoPilot/Agentic Tool Integration 5

3. From Model to System: How Modern AI Works 11

4. Actuarial Skill 17

5. From Tasks to Chains: Optimizing Workflow Efficiency 19

6. New AI Risks and Information Overflow 21

7. AI and Climate 24

Introduction

This addendum to the CAS AI Primer: Practical Guidance for Actuaries, Second Edition provides deep dives into key artificial intelligence (AI) topics based on the CAS AI Primer Working Group's research and discussions. They explore how AI impacts the actuarial profession and offer insights to help shape our future. We update these topics regularly and welcome your feedback and suggestions for new topics.

1. **AI Regulation and Its Impact on Actuarial Practice**: Analyzes evolving international AI regulatory mandates and their direct implications for the insurance industry and actuarial roles.
2. **Excel CoPilot/Agentic Tool Integration**: Demonstrates a loss triangle analysis using an Excel AI add-on.
3. **From Model to System: How Modern AI Works**: Discusses technical details of where the knowledge is stored, technologies to transform large language models (LLMs) into agents, and multi-agent workflows.
4. **Actuarial Skill**: Discusses what the actuarial profession and role could look like in the AI-driven environment of the future.
5. **From Tasks to Chains: Optimizing Workflow Efficiency**: Evaluates framework decision-making for automation, determining where full automation suffices versus where AI adds distinct value.
6. **New AI Risks and Information Overflow**: Discusses new risks and data overload introduced by AI.
7. **AI and Climate**: Examines data center energy consumption, environmental footprints, and their impact on the insurance industry.

1. US AI Regulation and Its Impact on Actuarial Practice

The US currently lacks a comprehensive federal framework governing AI in insurance. Instead, property and casualty (P&C) insurers navigate a complex patchwork of state-level regulations and guidance. A robust governance framework identifying and mitigating model risks — particularly concerning regulatory compliance, bias mitigation, and clear accountability — is imperative for successful AI adoption.

The NAIC Model Bulletin

Adopted in December 2023, the National Association of Insurance Commissioners (NAIC) Model Bulletin provides foundational guidance on the use of AI systems by insurers. Rather than creating new statutes, it outlines how existing unfair trade and antidiscrimination laws apply to AI. It mandates that insurers establish a written AI systems program that details governance, risk management, internal controls, and third-party oversight across the model life cycle. Crucially, this principles-based approach has gained significant regulatory traction; to date, approximately 25 states have adopted the bulletin, establishing it as the baseline standard for AI governance in the US insurance market.

Key state legislation

Several states have implemented more prescriptive regulatory requirements that directly impact P&C underwriting and pricing. For each framework, the burden on actuaries is shifting toward demonstrable fairness and transparency:

- **Colorado SB 21-169.** This law represents the strictest regime for algorithmic bias. It requires insurers to establish a risk-based governance framework and perform rigorous quantitative testing to ensure external data and predictive models do not unfairly discriminate against protected classes.
    - *Actuarial implications and key takeaways:* Actuaries must proactively develop and implement quantitative bias testing methodologies (e.g., disparate impact metrics) within their standard pricing and underwriting workflows. Model validation now explicitly requires proving that variables do not serve as proxies for race, gender, or other protected classes.
- **New York DFS Circular Letter No. 7 (2024).** Requires a documented, multistep fairness analysis before insurers rely on AI or external consumer data in underwriting or pricing, enforcing a strict proxy-testing mandate to prevent adverse disparate impacts.
    - *Actuarial implications and key takeaways:* Actuarial teams must integrate proxy testing directly into their model development pipelines before deployment. This shifts the burden of proof onto the actuary to document a verifiable fairness analysis for any external data source utilized.

- **Connecticut Bulletin No. MC-25 (2024).** Adopts the NAIC framework but introduces a state-specific annual "artificial intelligence certification" to ensure ongoing compliance with antidiscrimination laws.
    - *Actuarial implications and key takeaways:* Actuaries will play a central role in facilitating this annual certification. This necessitates building continuous monitoring systems and maintaining highly auditable compliance artifacts to quickly evidence model fairness on a recurring basis.
- **California CPPA ADMT Rules.** Regulations governing automated decision-making technology (ADMT) grant consumers transparency rights and the ability to opt out of AI-driven decisions that significantly affect their insurance coverage.
    - *Actuarial implications and key takeaways:* Models can no longer function as pure black boxes. Actuaries must design systems that not only accommodate human overrides but also produce plain-language explanations of how specific variables influence consumer outcomes to meet transparency and opt-out requirements.

International regulations

Because insurance models and vendor tools frequently cross borders, actuaries must also navigate a tightening global regulatory environment:

- **Canada: OSFI Guideline E-23 – Model Risk Management (2027)**
  Set to take effect in May 2027, the Office of the Superintendent of Financial Institutions (OSFI) has updated its model risk management guideline to explicitly cover AI and machine learning models at all federally regulated financial institutions, including insurers. It mandates a rigorous five-stage life cycle governance framework (design, review, deployment, monitoring, and decommission) and crucially extends model risk requirements to third-party vendor tools.
    - *Actuarial implications:* Actuaries must treat advanced AI and third-party "black box" vendor tools with the same rigorous governance applied to internal capital models. This includes creating robust explainability plans, checking training data for bias, and ensuring continuous independent reviews.
- **European Union: EU AI Act**
  The EU AI Act introduces a strict, risk-based classification system for AI. Under this framework, AI systems used for risk assessment and pricing in life and health insurance are explicitly classified as "high-risk" applications.
    - *Actuarial implications:* Actuaries operating in or intersecting with the EU market must establish robust explainable AI frameworks. The Act demands strict technical documentation, continuous human oversight, and data governance practices to ensure fairness. Furthermore, actuaries should note that even traditional complex actuarial models may fall under the EU's broad definition of AI, subjecting them to these stringent auditability requirements.

Practical implications for actuaries

These regulations reshape actuarial modeling processes, demanding rigorous compliance and standardized processes to ensure consistency and auditability throughout the deployment life cycle.

- **Professional accountability.** Under CAS and American Academy of Actuaries standards, the signing actuary remains accountable for the work product regardless of the tools used to produce it.
- **Methodology and auditability.** Actuaries must codify workflows using explicit numbered rules, factor selection logic, and escalation thresholds. All prompts, tool calls, and AI outputs must be retained and filed on an actuarial workpaper schedule consistent with professional standards.
- **Validation against "black boxes."** Since GenAI can function as a "black box," actuaries must apply well-defined evaluation criteria aligned with corporate risk appetite to assess both input data and model outputs.
- **Third-party oversight.** Organizations often rely on third-party foundation models. However, under these regulatory frameworks, the insurer retains ultimate compliance responsibility, necessitating careful vetting and monitoring.

Closing thoughts

The regulatory landscape for AI in insurance is rapidly transitioning from theoretical principles to enforceable, prescriptive mandates. We are seeing a distinct trend toward requiring quantitative proof of fairness (as pioneered by Colorado and New York) and holding institutions strictly liable for the logic inside third-party vendor tools (as emphasized by OSFI E-23).

Moving forward, actuaries should closely monitor the harmonization or divergence of state-level adoptions of the NAIC Model Bulletin, the potential expansion of the EU AI Act's "high-risk" classification into broader P&C lines, and the evolving technical standards for explainable AI. Ultimately, the actuarial profession is uniquely positioned to bridge the gap between technical model development and legal compliance, ensuring that AI adoption remains both deeply innovative and fundamentally ethical.

2. Excel CoPilot/Agentic Tool Integration

Shift. Until recently AI assistants could only talk to you about the data you gave them or that they could see, for example in Excel files. In 2026, a host of agentic tools have been released that more fully integrate work applications and AI model capabilities. Microsoft 365 Copilot and Anthropic's Claud for Excel can read multi-tab workbooks, write formulas and create pivot tables, build charts, and check their own work.

Use cases. The biggest win for this change in AI workflows, with respect to actuarial work, is the ability to synthesize large and complex Excel files. These models are very effective in helping you understand an Excel model with which you are unfamiliar, creating first-pass edits and reconciliation, as well as creating documentation for models.

Best practices. Vendors are increasingly routing users to cheap, fast models by default ("auto"). These are fine for very simple queries like creating a simple formula in a single worksheet. For complex tasks, use the models with the strongest reasoning capabilities (Opus-class in Claude, 5.6 Sol in GPT, and similarly for other labs). As always, regardless of the model chosen, there are always risks that the AI output includes mistakes. All output should be treated as a draft and reviewed thoroughly.

Going forward. Over time, the tools are shifting from a search box to a junior analyst, with integration of these models with the general office workflow happening in the future. Vendors will increasingly push users toward cheaper models due to near-term compute constraints.

Example for M365 Copilot

Relevant files

Example video (gif — lower quality): https://imgur.com/a/mwE3dix

Mack method original file: https://www.casact.org/sites/default/files/2021-02/education_spring_2006_handouts_mack-method-handout.xls

Mack method (AI edited): https://docs.google.com/spreadsheets/d/1_tKfGYVXuRqrG9YgPr_QHR8hy-LuJ5cI/edit?usp=drivesdk&ouid=116042963370858028545&rtpof=true&sd=true

M365 Copilot text outputs

Prompt 1:

Can you please explain to me what is happening in this Excel workbook? Thank you!

Response 1:

I've traced every formula in the workbook. Here's what it's doing.

Overview

This is an actuarial **loss reserving** model implementing Thomas Mack's "Distribution-Free" method (1993) — the industry-standard way to put a *standard error* (uncertainty measure) around a chain

ladder reserve estimate. It uses Mack's famous published example dataset (the "Taylor and Ashe" triangle), and column X checks the results against Mack's published answers.

Everything lives on one sheet, Mack Method, in A1:X98.

Section-by-section walkthrough

1. The loss triangle — Mack Method!A4:K15

Rows are accident years i = 1–10; columns are development ages k = 1–10. Values are **cumulative paid losses**. It is triangular because older years have more development history — year 1 is fully developed to age 10 ($3,901,463), year 10 only has age 1 ($344,014).

2. Age-to-age factors — Mack Method!A18:J30

Each cell is simply the ratio of one period to the prior, e.g., Mack Method!B19:

=C6/B6

Mack Method!B29 is the volume-weighted average development factor (Mack's $\hat{f}_k$):

=SUM(C6:C14)/SUM(B6:B14)

Mack Method!B30 chains these backward into cumulative **loss development factors (LDFs)**: =B29*C30, giving 14.447 at age 1 down to 1.000 at age 10.

3. Reserve projection — Mack Method!M6:P17

- M = latest **diagonal** (most recent cumulative loss per year)
- N = the applicable **LDF**
- P **Ultimate** = =M6*N6
- O **Reserve** = =P6-M6 (money still to be paid)

Total reserve: $18,680,856 on $34,358,090 paid to date.

4. Variance estimation — Mack Method!A35:J47

Mack Method!B36 computes Mack's squared residuals, weighted by volume:

=B6*(B19-B$29)^2

Mack Method!B46 turns those into Mack's $\hat{\sigma}_k^2$ estimator:

=SUM(B36:B44)/(COUNTIF(B36:B44,">0")-1)

Mack Method!J47 is the classic **tail extrapolation** for the last column (where there are no data to estimate variance):

=I47*MIN(H47,I47)/MAX(H47,I47)

5. The two variance multipliers — Mack Method!A50:J56

Mack splits uncertainty into two pieces:

	Formula	Meaning
Process variance Mack Method!B51	=B30*B47/B29^2	Pure random chance in future claim payments
Parameter variance Mack Method!B55	=(B$47/B$29^2)/SUM(B$6:B$14)	Uncertainty in the estimated development factors themselves

Rows 52 and 56 accumulate these recursively across all remaining development periods (=B51+C52).

6. Results — Mack Method!Q6:X17

- R **Process SD** = =(P7*Q7)^0.5
- T **Parameter SD** pulls from the correlation matrix in Mack Method!B90:B98
- V **Total SD** = =(R7^2+T7^2)^0.5 (Pythagorean combination)
- W **CV** = coefficient of variation = SD ÷ Reserve

The **total row 17** is the important part. Process error aggregates independently (=SUMSQ(R7:R15)^0.5), but parameter error is **correlated across accident years**, so it must be summed over the full covariance matrix:

=SUM(C77:K85)^0.5

7. Matrices — Mack Method!A61:K98

These three 9×9 grids exist purely to handle that correlation:

- Mack Method!B62:K72 — **estimation error factors**, using =IF($A64<C$62,$B64,C$63) to pick the shared (minimum) development history between each pair of years
- Mack Method!B75:K85 — the **covariance matrix**, factor × Ultimate*i* × *Ultimate*j
- Mack Method!B88:K98 — the same, normalized into **correlation coefficients**

Key results

Accident Year	Reserve	Total SD	CV
2	$94,634	$75,535	80%
3	$469,511	$121,699	26%

5	$984,889	$261,406	27%
7	$2,177,641	$558,317	26%
9	$4,278,972	$971,258	23%
10	$4,625,811	$1,363,155	29%
Total	**$18,680,856**	**$2,447,095**	**13%**

The headline: the reserve estimate is **$18.68 million ± $2.45 million**, a coefficient of variation of **13%**.

Two things worth noting:

1. **Diversification works.** Individual years have CVs of 19–80%, but the total is only 13% — errors partially offset across years.
2. **The model validates perfectly.** Every value in column X (Mack's published answers) matches column W exactly — 80%, 26%, 19%, 27%, 29%, 26%, 22%, 23%, 29%, and 13% total. This workbook is a verified replication of the textbook example.

Would you like me to adapt this to your own loss triangle, or add anything like a chart of the reserve ranges?

Prompt 2:

Can you please create a new sheet where I have an interactive assumptions panel? I want to be able to select the development factor/the development factor methodology, tail factor, and confidence level. Please format the input and output cells differently, so I can easily differentiate between them. Thank you!

Response 2:

The panel validates perfectly — with volume-weighted and a 1.000 tail it reproduces the original model's **$18,680,856** reserve and **$2,447,095** SD exactly, with a $0 variance versus baseline.

Assumptions sheet created

A fully interactive control panel, placed as the first sheet. Everything is formula-driven and linked live to Mack method — no hard-coded numbers.

Color coding

Style	Meaning
🔵 **Blue fill, blue bold text, thick border**	**INPUT** — you edit these

⬜ **Gray fill**	**OUTPUT** — formulas, don't type here
🟡 **Yellow fill**	Derived helper (z-score)
🟢 **Green fill**	Key selected results
🔷 **Dark blue**	Section headers / totals

A legend sits at Assumptions!F5:J5.

The four inputs — Assumptions!C7:C10

1. **Development factor method** (Assumptions!C7) — dropdown with five options: Volume-Weighted Average, Simple Average, Latest 5, Latest 3, Medial (ex hi-lo)
2. **Tail factor** (Assumptions!C8) — numeric, validated 0.5–3.0
3. **Confidence level** (Assumptions!C9) — dropdown: 50%, 75%, 80%, 90%, 95%, 97.5%, 99%, 99.5%
4. **Distribution** (Assumptions!C10) — Lognormal or Normal

Each has a hover tooltip explaining its use.

How it works

All five methodologies compute simultaneously in Assumptions!C17:L21. The selected one is pulled through by Assumptions!C22:

=INDEX(C17:C21,MATCH(C7,{"Volume-Weighted Average";"Simple Average";"Latest 5";"Latest 3";"Medial (ex hi-lo)"},0))

Your tail factor lands at age 10, then LDFs chain backward in Assumptions!C23:

=C22*D23

The lognormal confidence bound at Assumptions!J39 correctly handles reserve skewness:

=IF(C10="Lognormal",F39*EXP(C12*SQRT(LN(1+H39^2))-0.5*LN(1+H39^2)),F39+C12*G39)

Current output (defaults: Volume-Weighted, tail 1.000, 95%)

Metric	Value
Total Reserve	**$18,680,856**
Total SD	**$2,447,095**
CV	**13.1%**

Reserve at 95%	**$22,955,180**
Risk Margin	**$4,274,324**
Change vs. Baseline	**$0** ✓

Verified sensitivity

I tested three scenarios and all flowed through correctly:

- **Latest 3 + 1.05 tail + 99%** → reserve rises to $20,650,255 (+10.5%), 99th percentile $27,734,158
- **Switching to Normal** at the same settings → $26,943,202, about $791k lower than lognormal, which is the expected behavior since the lognormal's right tail is fatter

References

Mack, Thomas. 1993 "Distribution-Free Calculation of the Standard Error of Chain Ladder Reserve Estimates." *ASTIN Bulletin* 23: 213–25. https://doi.org/10.2143/ast.23.2.2005092.

3. From Model to System: How Modern AI Works

An intro to the terminology and architecture behind current AI systems, prepared for actuaries.

Earlier guidance in the CAS AI Primer addressed how to prompt a model and validate its output. Attention in the field has shifted from how to prompt a model and validate its content to the system built around the model: the agent, the harness, and the protocols that connect them. This section traces why each of these pieces emerged and what problem it solves and connects each one to actuarial work.

> **The Organizing Idea**
>
> A language model operates under three constraints: it retains nothing between requests, cannot act on its own, and knows only the text placed directly before it. Every term that follows describes something built to work around one of these three constraints, and understanding why each one appeared, not only what it is, is the fastest way to make sense of the terms.

What is the model: Where its knowledge lives

A model's knowledge can be built in two different ways. One approach changes the model itself before it is ever used: **fine-tuning** retrains it on an organization's own data, and a newer, less costly method called **RLVR** rewards it for producing answers that pass an automated check, embedding the knowledge directly into its parameters. The result is permanent and private, but expensive to revise. The alternative is to supply what the model needs at the moment of use through context engineering, rather than building it in beforehand. A parallel decision has become live as the built-in option matured: whether to run an **open-weight** model, hosted on an organization's own infrastructure, or a **closed-weight** model, reached only through an external API. By mid-2026, open-weight models had grown capable enough that hosting sensitive data entirely in-house was a realistic choice rather than a compromise.

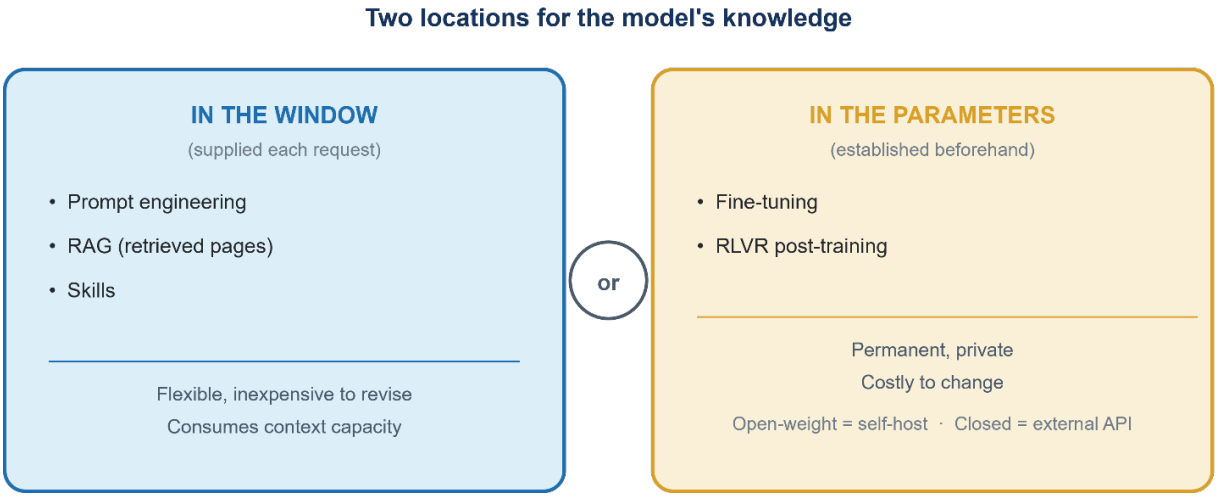


Figure 1. Knowledge may sit in the context window or be built into the model's parameters.

ACTUARIAL EXAMPLE

Travelers took this approach with TravelersLLM, trained on internal documents to outperform general models on insurance tasks at lower cost, an example of the model-ownership strategy insurers are increasingly pursuing.

Context engineering: What the model sees

Early attempts to get useful work from a model relied almost entirely on writing a good prompt, a single well-crafted instruction placed inside what is called the **context window**, the fixed span of text a model can read at once. As tasks grew more complex, a second technique took hold: **RAG**, retrieval-augmented generation, which fetches only the passages relevant to a question rather than handing over an entire document. A third followed as procedures multiplied: **skills**, which keep detailed instructions outside the window and load them only when a task calls for them. What these techniques share became clear once teams started overloading the window instead of curating it: recall degrades as the window fills, a problem now called **context rot**, and the discipline of avoiding it, deciding precisely what a model sees and when, is what practitioners call **context engineering**.

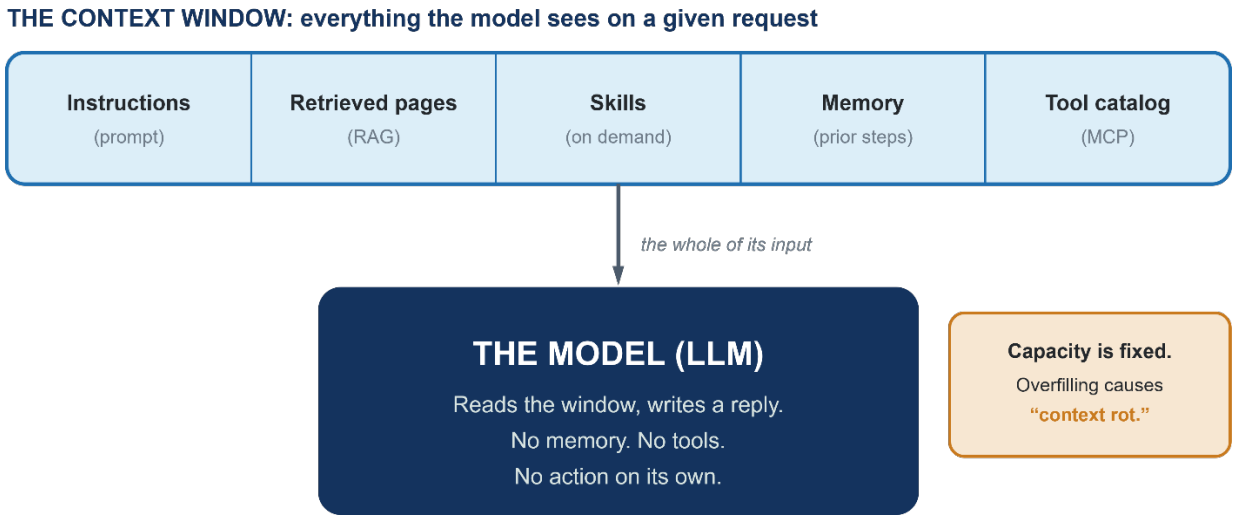


Figure 2. The context window is the model's entire input, and its capacity is fixed.

ACTUARIAL EXAMPLE

Summarizing a 200-page claim file in full risks context rot and a missed detail. Retrieving only the loss description and coverage clause, with the reserving methodology held as a skill, produces a sharper result from the same model.

The harness: How the work gets done

A model can state that a loss triangle should be retrieved, but on its own it cannot retrieve anything: it only produces text. That execution is performed by the **harness**, software products of many different kinds, built around a model to give it those abilities. For example, Claude Code and Cursor are harnesses built for coding work; in insurance, the same idea shows up in platforms such as Sixfold and Kalepa. Each wraps a model with the tools and memory underwriting work depends on. Whatever the specific product, a harness assembles the context, carries out the model's requests, returns the results, and repeats the cycle. It also maintains **memory** between steps, manages failures, and enforces the points at which a person must review the work.

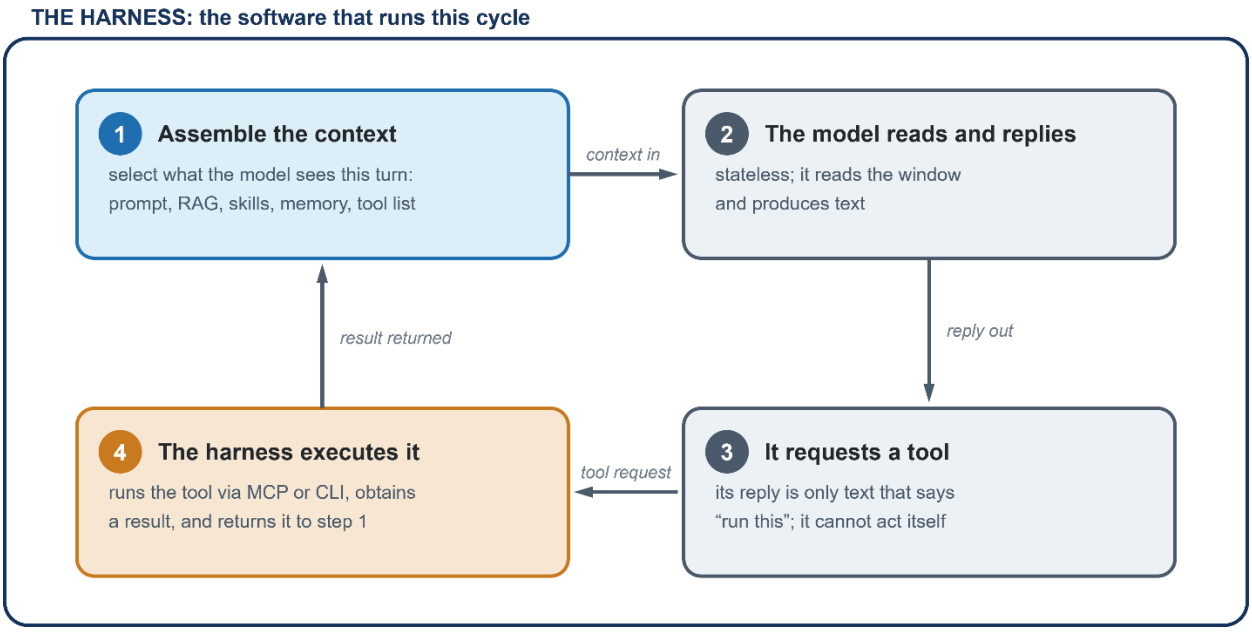


Figure 3. The model performs only step 2; the harness performs steps 1, 3, and 4.

> **ACTUARIAL EXAMPLE**
>
> In a quarterly reserve review, the harness pulls loss triangles from the claims system, retries automatically when a data query times out, recalls the prior quarter's selected assumptions when the model references them several steps later, and halts the process, requiring a human actuary to review and approve the figure before it is finalized. The model selects and interprets the factors; the harness handles retrieval, continuity, and control; the actuary retains final judgment.

Connecting to tools and other agents

MCP emerged in late 2024 as the standard way to connect a model to external tools and data, and companies building AI systems in-house adopted it quickly, largely because it logs and authenticates every action, producing exactly the audit trail governance and compliance teams ask for. As agent use grew heavier through 2026, a criticism took hold: MCP's tool definitions must sit inside the context

window, and by some measures this consumes several times more capacity than a simpler alternative. That alternative, **CLI**, lets the model issue commands directly, with nothing to preload, which is far more efficient but leaves no inherent record of what was done. The resulting debate, over whether MCP was being displaced, settled in practice on neither side winning outright: CLI suits lean, uncontrolled work, while MCP remains preferred wherever an organization must be able to show what its AI system did and why. A newer protocol, **A2A**, introduced in 2025, addresses a different problem: coordination between agents built by different organizations, which has become more relevant as multiagent systems, rather than single agents, become the norm.

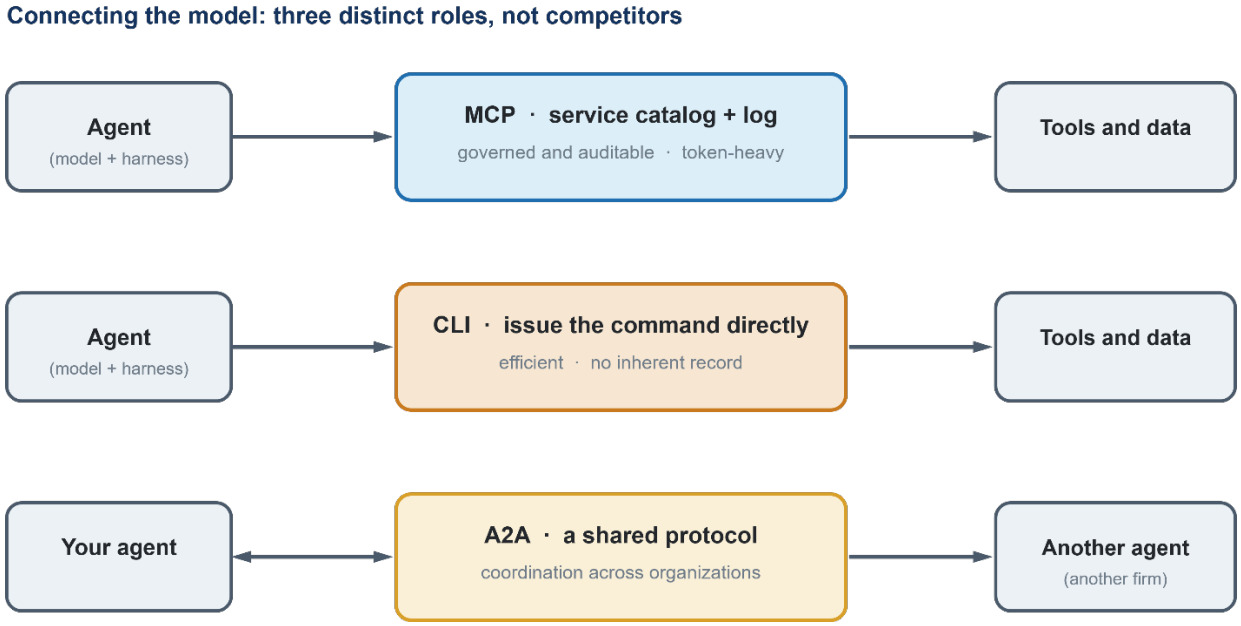


Figure 4. MCP, CLI, and A2A perform three separate roles rather than competing for one.

ACTUARIAL EXAMPLE

For regulated work, MCP's built-in logging gives insurers an audit trail they can show to risk teams or regulators since it is easier to produce the kind of documentation the NAIC Model Bulletin requires insurers to maintain. For lower-risk, internal tasks, CLI's speed is often the better fit. The right choice depends on the setting, not a fixed rule.

Scaling up: Agents and orchestration

Put a model and a harness together in a loop that runs until a task is complete, and the result is what the field calls an **agent**. A single agent handles a contained task well, but as organizations pushed agents toward larger jobs, a familiar problem returned: one agent's context window fills up and its performance degrades in the same way an overloaded prompt does. The response, **multiagent orchestration** (sometimes called **graph engineering**), splits the job across several agents, each working a narrow, verifiable slice and reporting a short summary back to a lead agent rather than

passing along everything it did. The rule that has emerged for dividing the work is **contract-first decomposition**: a subtask is handed off only when what counts as a correct result can be defined before the work begins.

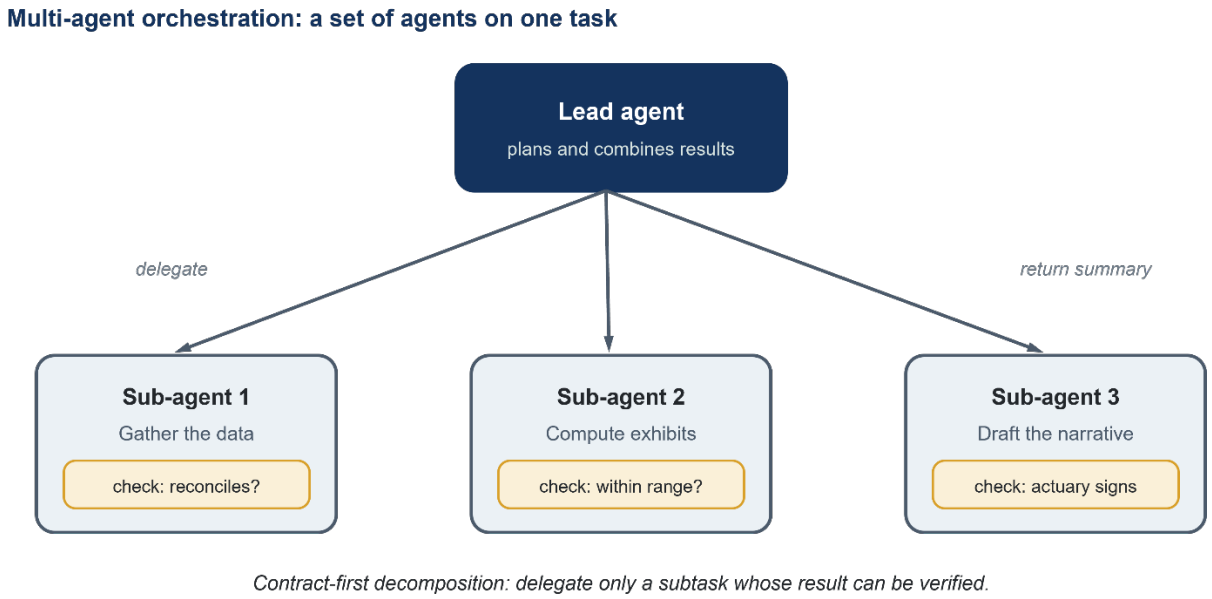


Figure 5. Each agent is responsible for a verifiable part and reports to a lead agent.

> **RELATION TO ACTUARIAL PRACTICE**
>
> This mirrors the actuarial control cycle: decompose the task, verify each step, and keep a qualified person accountable at the checkpoints. In a statutory filing, one agent gathers data subject to a reconciliation check, another computes exhibits subject to a range check, and a third drafts the narrative for the appointed actuary's sign-off.

How this is developing, and what it means for us

These methods did not appear at once; they arrived in stages, prompting and retrieval first, then the protocols that connect a model to tools, then the harness and context engineering built around it, and most recently the coordination of multiple agents. None of this has settled. The pace of change has accelerated, and the terms in this section should be expected to keep shifting as the underlying technology does.

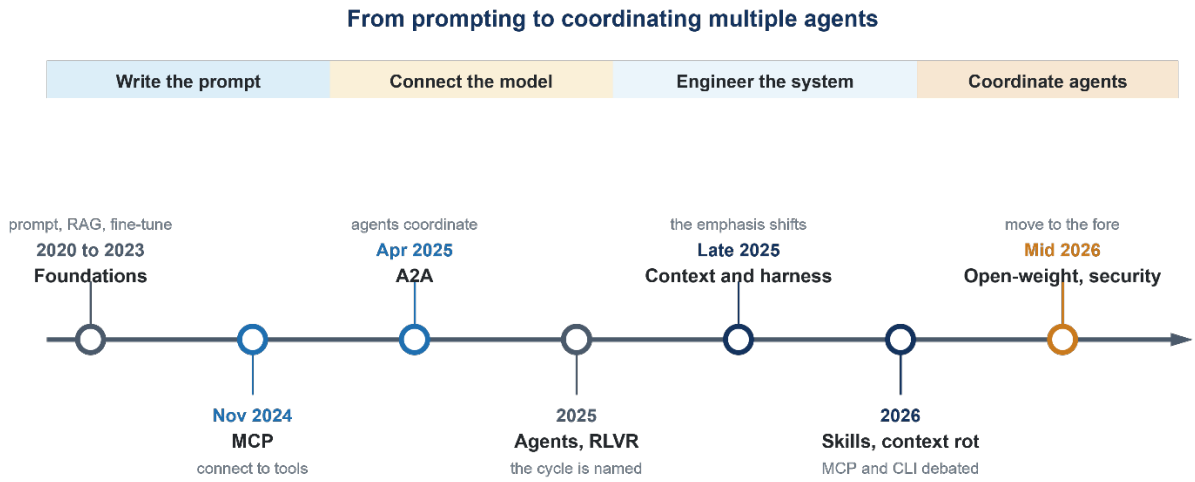


Figure 6. Four stages: prompting, connection, system engineering, coordination. Note: Model names and dates reflect information available as of mid-2026.

For actuaries, the practical takeaway is simple: The vocabulary and tools in this section will keep changing; staying current matters as much as understanding today's version of them. Trust, security, and verification, prompt injection, memory poisoning, and evaluation pitfalls like reward hacking, are covered in the AI Risk section.

4. Actuarial Skill

Actuaries and data scientists occupy a unique position in the insurance industry, possessing the technical expertise required to locate, manipulate, and structure complex data.

While advancements in AI streamline manual actuarial tasks, they simultaneously unlock numerous opportunities for higher-value, impactful work that was previously unattainable — the true source of value generation.

The flowchart that follow displays two paths into the future. One future leads to fewer actuaries and the other leads to more. We believe we are on the “more actuaries” path due to the competitive pressures that AI enables. Someone must make sense of what the AIs do, and that someone needs years of highly focused, specialized training and experience — like an actuary.

More actuaries or fewer?

More

New creative solutions to existing problems as actuaries utilize AI to do work they couldn't before

New risks emerge for actuaries to quantify as competition drives the industry to do the new things

Fewer

Manual cyclical projects turned over to automations

AI explanations of trends and data given full credibility

Examples of new work

- Measure AI workflow accuracy
- Monitor AI performance on an ongoing basis
- Analyze claims using data enriched by AI
- Enhance pricing models with LLM-derived data
- Evaluate IBNR adequacy with broader considerations
- Design AI-enabled workflows to improve overall efficiency
- Design monitoring systems with live data ingestion and immediate alerts on anomalies
- Rethink risk management for insurers in light of the new norm of human-AI interaction.

New skills

- Prompt engineering
- Goal and specification authoring
- Version control systems
- Collaborative programming
- Managing collaborative programming
- Live data ingestion
- System engineering

5. From Tasks to Chains: Optimizing Workflow Efficiency

Modern insurance companies face immense pressure to optimize workflows, accelerate speed-to-market, and manage rising risk volatility. Incorporating AI into their systems for enhanced efficiency is often seen as an "easy" solution to reduce costs and improve performance.

This section explores whether AI is the appropriate tool for driving these improvements and outlines the key considerations necessary to achieve meaningful process transformation.

Deterministic rules versus probabilistic adaptation

Traditional automation operates on a predefined set of rules or workflows. It is the classic "if this, then that" of system design. Typical examples include running database scripts, segmenting drivers into age buckets, and calculating various loss development factors (LDFs) for reserving analysis. Traditional automation is highly reliable, predictable, and computationally efficient. It is entirely rigid; it can only do what it is explicitly programmed to do. If raw data schemas change, or an unexpected scenario arises, traditional automation halts and requires manual human intervention

In contrast, AI is designed to learn, adapt, and make decisions based on data. Rather than following a fixed, coded script, AI systems analyze complex multivariable patterns, recognize anomalies, and improve their accuracy over time. AI is probabilistic rather than deterministic. This means it is built to handle ambiguity, respond to changing contexts, and make predictions where explicit mathematical formulas or hard rules are impossible to write

Where to start: Write out the rules first

Determining whether to use automation or AI begins by defining a task's rules. Processes with explicit, repeatable steps are best handled via code. Conversely, tasks involving nuance — such as classifying free-text loss descriptions — often resist rigid, rule-based automation due to infinite variations in reporting style. This evaluation is inherently valuable: A complete rule set forms an automation specification, while an impasse identifies precisely where human judgment — and thus AI — is necessary.

Most actuarial workflows, such as quarterly reserve reviews (data extraction, reconciliation, triangle building, and factor computation) and rate impact analyses, are routine and well suited for traditional automation. These processes prioritize consistency, predictability, and auditability — attributes where deterministic scripts excel.

However, key judgmental components remain. In reserving analyses, actuaries must synthesize complex variables like macroeconomic shifts, geopolitical tensions, and evolving claim-handling practices to select loss development factors. When historical data are insufficient — such as during a pandemic — AI can be incorporated to synthesize diverse internal and external inputs, offering recommendations that augment human expertise. Decomposing these workflows is essential; it isolates tasks requiring genuine judgment from the automated bulk, guiding smarter development.

Finally, AI is invaluable even when a process is better automated. It excels at translating methodology into code, debugging, and explaining technical logic, allowing actuaries with deep domain expertise to design reliable, reusable scripts without AI becoming part of the runtime production process.

Node-level efficiency versus system-level efficiency

Rushing into AI implementation without first auditing underlying data quality and IT infrastructure remains a critical pitfall. Workflow bottlenecks frequently stem from legacy systems and manual data handoffs rather than a lack of advanced AI capabilities. Consequently, integrating AI piecemeal into isolated operational steps creates long-term inefficiencies by requiring continuous human intervention to bridge structural gaps.

To move beyond incremental productivity gains and realize true organizational transformation, actuarial processes must be restructured around **system-level efficiency** rather than isolated **node-level efficiency** (Demirer et al. 2026). Targeting AI applications improve specific tasks like report generation, data validation, etc., and leave the sequential workflow intact. As these optimized nodes remain siloed by manual handoffs, overall throughput is still fundamentally constrained by human processing speed.

To overcome these limitations, research from MIT Sloan highlights the strategy of **task chaining** — connecting adjacent, compatible tasks into automated, end-to-end pipelines. By executing calculations and output generation seamlessly across steps, task chaining eliminates the friction and delay caused by manual data transfers (Demirer et al. 2026).

While maintaining "human-in-the-loop" is essential, requiring manual verification at every transition point introduces significant friction. Frequent context switching disrupts focus, leads to superficial approvals, and increases operational costs — ultimately eroding the efficiency gains provided by automated systems.

To minimize AI-human collaboration cost, organizations must critically reevaluate when and how human intervention occurs. The core objective should be engaging human expertise only for critical sign-offs or when predefined exception thresholds are triggered. Rather than eliminating human involvement, this approach ensures human judgment is strategically applied where it adds the highest value, freeing actuaries to concentrate on strategic decision-making.

References

Shahidi, P., Demirer, M., Horton, J. J., Immorlica, N., & Lucier, B. (2026). *Chaining Tasks, Redefining Work: A Theory of AI Automation*. MIT Sloan School of Management / Microsoft Research.

6. New AI Risks and Information Overflow

AI has introduced novel operational and systemic risks across the insurance industry. In this section, we examine two quantifiable examples of AI-driven risk that insurers now encounter, along with specific metrics actuaries can use to evaluate them. Additionally, we address the challenge of information overflow, exploring how AI-enhanced productivity and content generation impact actuarial workflows and decision-making.

Agent-driven erroneous decisions at machine speed

For an insured business, an agent could place excessive orders, expose confidential information, deploy defective code, or make unauthorized contractual representations. The resulting losses may trigger several types of coverage, including cyber, crime, errors and omissions, directors and officers, product liability, and business interruption.

Actuaries should begin measuring:

- The number of agents possessing write, payment, deployment, or contractual authority
- Transaction volume and maximum financial authority per agent
- The number and criticality of systems accessible to each agent
- The proportion of actions requiring independent human approval
- The number of downstream automated actions triggered by one initial decision
- Detection, containment, and rollback times
- Error and near-miss rates observed in testing and production

Foundation model use as concentration risk

Similar to AWS data centers representing a concentration of risk in processing power, the AI models represent a concentration of decision patterns used across entire industries. A common model may generate plausible but systematically wrong outputs. The resulting event is therefore not necessarily a visible system failure; it may be a prolonged accumulation of defective decisions.

A single upstream failure could generate simultaneous claims under technology errors and omissions, cyber, professional liability, employment practices, product liability, business interruption, and other policies. Diversification by geography or industry may provide little protection when the common dependency is technological.

Actuaries should measure:

- The percentage of insured limits dependent on each model provider and model family
- Direct dependencies and less-visible dependencies embedded in third-party software
- Concentration by model version, hosting provider, and training lineage
- Expected correlation among failures involving the same model

- Time required for policyholders to substitute another provider or revert to manual processes
- Probable maximum loss from a one-day, one-week, or one-month common-model failure
- Potential overlap among affected policies and coverage lines

Information overflow

In addition to the risks related to systems and models, the rapid advancement of AI is fundamentally compressing traditional insurance project cycles, accelerating everything from research and rate filings to IBNR decisions. Simultaneously, robust AI search and summarization tools are inundating reviewers with unprecedented volumes of information. Consequently, this dual pressure — increased task velocity and information saturation — precipitates severe cognitive fatigue. Under the urgency to deliver results at record speeds, reviewers risk approving outputs that merely "look good," often bypassing the rigorous downstream impact analysis essential to protecting policyholders and industry stability. Alternatively, this cognitive saturation forces organizations to deploy secondary filtering AIs to monitor and summarize primary AI outputs. Ultimately, this dynamic establishes a closed loop of absolute technological reliance. As humans are pushed further up the abstraction ladder, they risk losing the ability to manually audit underlying processes, thereby entrenching systemic operational vulnerabilities.

To effectively manage information overflow, actuaries must distinguish between value-generative initiatives and value destructive AI implementation.

Consider reserving analysis: Value-generative AI acts as an insight synthesizer, integrating qualitative perspectives across reserving, claims, and pricing to inform IBNR selection. Conversely, a non-value-adding implementation might directly insert an AI model into existing workflows to select loss development factors based solely on historical patterns. In that case, automation removes human judgment from detailed analysis while adding a supervisory review burden, delivering no incremental value beyond basic pattern matching. Similarly, in pricing, value-generative AI synthesizes insights from previously hard-to-consume resources — such as regulatory publications, industry reports, and competitor analysis — to enrich current evaluations. By contrast, a non-value-generative initiative might build a chat-based research assistant for pricing and loss analysis within existing data structures. Such a tool can obscure underlying data, prioritizing convenience over data fluency and enabling analysts to bypass direct data manipulation.

These situations can be effectively mitigated by establishing methods to thoroughly quantify the financial and economic costs and benefits of each initiative:

- **Financial costs and benefit.** What are the comprehensive implementation costs, including software, ongoing training, and maintenance and how much benefit can be quantified?
- **Nonfinancial costs and benefits.** How much time and effort can be reclaimed, and how effectively can those saved resources be redirected toward higher-value initiatives? Furthermore, what is the associated collaboration cost or hidden operational burden, such as increased review time, system overreliance, or workforce fatigue?

- **Innovation versus automation.** Are we leveraging AI to enable new capabilities, or are we simply using it to automate existing problems without improving our methodology?
- **Strategic alternatives.** Have we adequately explored alternative solutions, including non-AI approaches?
- **Problem definition.** What fundamental issue are we solving, and what is the concrete, quantifiable benefit to the company and the society that AI provides?

Future market leadership will not be defined by those who chase the latest trends, but by those who can translate AI initiatives into systematic enhancements that materially improve operational efficiency and long-term value.

7. AI and Climate

The case of the data center

Public opposition to AI data centers is surging, with seven in ten Americans opposing construction in their local areas (Jones 2026). Politicians are rapidly responding: Pennsylvania enacted blocking authority for new facilities, while Texas established strict water, electricity, and quality-of-life standards — illustrating a growing divide between outright bans and negotiated conditions.

It is worth noting that global data-center electricity consumption is projected to reach 945 TWh by 2030 (up from 415 TWh in 2024), with AI-accelerated servers growing nearly 30% per year and the US and China accounting for roughly 80% of global demand growth through 2030 (International Energy Agency 2025). More recent data suggest the pace is accelerating rather than leveling off: Electricity consumption from AI-focused data centers surged 50% in 2025 alone, and capital expenditure by the largest technology companies exceeded $400 billion that year, with another 75% increase expected in 2026 (International Energy Agency 2026). In the US specifically, data-center power demand is forecast to more than double, from 31 GW in 2025 to 66 GW in 2027, driven by the accelerating AI infrastructure buildout (Goldman Sachs 2026).

That growth is already showing up on utility bills. Researchers estimate that existing data centers have pushed average wholesale electricity prices 2% to 6% higher nationwide, with larger surges in the areas where facilities are most concentrated — a dynamic that helps explain why local opposition tends to cluster geographically rather than spreading evenly across the country.

Environmental impact is highly context dependent

The most visible environmental factor is electricity use, which translates into greenhouse gas emissions, depending on when and where that power is drawn. A data center on a coal- or gas-heavy grid carries a very different carbon burden than one running on low-carbon electricity.

One major consequence is urban heat. The latest science shows that the waste heat produced by data centers can increase surrounding land surface temperatures by as much as 16°F and increase air temperatures in nearby neighborhoods by up to 4°F. Given that cities are typically 1–7°F warmer than surrounding areas already, additional warming from data-center waste heat can demand more air conditioning energy and increase human health risks in nearby communities (Yale School of the Environment 2026).

Water and land use compound the picture. By 2030, the same data-center buildout projected to consume 945 TWh of electricity is expected to carry a water footprint of 9.3 trillion liters, equal to the basic annual domestic water needs of all 1.3 billion people in Sub-Saharan Africa, and a land footprint of roughly 14,500 km^2, about twice the size of the Jakarta metropolitan area. The same analysis projects 2.5 million metric tons of AI-related e-waste by 2030 and calls for mandatory environmental disclosure by AI providers (UNU-INWEH 2026).

Some counter arguments for data-center environmental concerns

- History has shown that increases in electricity demand tend to spark clean-energy booms. When the US needed lots of power after WWII, it built nuclear plants. When China needed power recently, solar and wind took off. Data centers could do the same thing now.
- AI companies are the perfect customers for new clean energy. They need power 24/7, they need a lot of it, and they can afford it because electricity is a small slice of what a data center costs compared with chips. Right now, they are using natural gas because it is the only reliable option, but they are also funding next-generation nuclear, geothermal, and carbon-capture gas plants, technologies that need deep-pocketed buyers to get off the ground (Nordhaus 2026).
- Water impact depends heavily on cooling technology, and the technology has improved substantially. Many newer facilities use closed-loop liquid cooling, where water is filled once and then recirculated continuously within sealed pipes rather than evaporated or discharged, leaving day-to-day water consumption comparable to, or even less than, a commercial office building of similar size. There is still a meaningful onetime water draw during construction (a medium-sized closed-loop facility can require water roughly equivalent to half an Olympic-sized pool), but after that the system is sealed and self-contained. Some operators, including CoreWeave, now build all new capacity to this standard by default.

The optimists are describing what could happen if the demand is channeled well; the opponents are describing what is happening in their counties right now. The two sides are not quite arguing about the same thing.

What this means for actuaries

The data-center buildout is not just an environmental or political story, it is fast becoming a new exposure class, and actuaries touch it from several directions at once — as pricing actuaries on the property, energy, and liability lines that now insure this infrastructure; as reserving actuaries watching for emerging loss patterns with almost no historical data behind them; and as enterprise risk actuaries assessing how concentrated, correlated this exposure is becoming within a book.

Pricing considerations

- *Grid- and power-interruption risk.* With US data-center demand set to more than double by 2027 (Goldman Sachs 2026) and wholesale electricity prices already up 2% to 6% nationally in concentrated areas, business interruption and contingent business interruption pricing for data-center-dependent operations, and for the grid customers around them, needs to account for a tightening, more volatile power market rather than historical utility-rate stability.
- *Water scarcity as a rating variable.* The UNU-INWEH (2026) water-footprint projections suggest that facilities sited on evaporative or open-loop cooling in water-stressed regions carry materially different risk than closed-loop facilities. Underwriting and

pricing questions actuaries should be asking cedants and insureds: which cooling technology is in use, what is the facility's water source, and is it located in a basin already flagged for stress? This is a rating factor that barely existed five years ago for most property books.

- *Property and construction risk.* Land footprint growth (roughly 14,500 km^2 of new development projected by 2030) implies a wave of new construction risk — course-of-construction, builder's risk, and fire/property lines, concentrated in a small number of counties, often ones without much prior experience insuring heavy industrial-scale electrical and cooling infrastructure.
- *Casualty and environmental liability.* Urban heat effects (up to 16°F land surface increase, 4°F ambient air increase per Yale School of the Environment 2026) and the projected 2.5 million metric tons of AI-related e-waste by 2030 (UNU-INWEH 2026) open the door to nuisance, toxic tort, and environmental liability claims that pricing actuaries have limited precedent for, closer in shape to early-stage PFAS or historical environmental liability than to standard CGL experience.

Reserving considerations

- *No credible loss history.* Because the exposure is only a few underwriting cycles old, reserving actuaries are working with thin or no triangle data. Loss development patterns will need to borrow from adjacent classes (heavy industrial power generation, telecom infrastructure, environmental liability) until the class matures, with wide margins for uncertainty.
- *Emerging and latent claims.* E-waste and heat-related health claims, in particular, may develop on a long tail similar to environmental and products liability; claims that are not yet visible in the data but that the exponential growth curves (12% → 15% → 30% annual growth in different International Energy Association scenarios) suggest are being underwritten today in growing volume.
- *Regulatory tail risk.* UNU-INWEH's call for mandatory environmental disclosure, alongside actions like Pennsylvania's blocking authority and Texas's negotiated standards, signals that the regulatory and disclosure regime around data centers is still being written. Reserves should build in some allowance for retroactive standards, retrofit mandates, or disclosure-driven litigation that could reshape loss patterns after policies are already on the books.

Aggregation, catastrophe, and reinsurance considerations

- *Geographic and grid concentration.* Because roughly 80% of global demand growth is concentrated in just the US and China (International Energy Association 2025), and because capital is flowing toward the same handful of grid regions and water basins, insurers and reinsurers building books in this space face real aggregation risk, a regional

grid failure, heat wave, or drought could simultaneously impact power interruption, property, and liability exposures across many insureds at once.

- *Correlation with climate catastrophe models.* Because data centers are disproportionately sited to chase cheap land and power rather than to avoid catastrophe-prone geography, cat modeling teams should check whether data-center concentrations are inadvertently building up in wildfire, drought, or extreme-heat zones already flagged by existing climate models — the risk stacks rather than diversifying.
- *Capacity and reinsurance appetite.* With capex from the largest technology companies exceeding $400 billion in 2025 and projected to grow another 75% in 2026 (International Energy Association 2026), insured values in this class are scaling far faster than actuaries' ability to build credible pricing models — a mismatch that reinsurers are likely to price cautiously — or ration capacity for — until loss experience matures.

The open question for the profession

Energy, AI, and their actuarial implications remain under researched relative to how fast the exposure is growing. The data-center buildout is a live case study in a broader pattern; infrastructure scaling faster than the actuarial data needed to price, reserve, and reinsure it responsibly.

References

Goldman Sachs Exchanges. 2026. *The Outlook for Data Center Power Demand as AI Token Use Grows*. https://www.goldmansachs.com/insights/goldman-sachs-exchanges/the-outlook-for-data-center-power-demand-as-ai-token-use-grows.

International Energy Agency. 2025. *Energy and AI*. IEA. https://www.iea.org/reports/energy-and-ai.

International Energy Agency. 2026. *Key Questions on Energy and AI*. IEA. https://www.iea.org/reports/key-questions-on-energy-and-ai.

Jones, Jeffery M. 2026. *Americans Oppose AI Data Centers in Their Area*. Gallup. https://news.gallup.com/poll/709772/americans-oppose-data-centers-area.aspx.

Nordhaus, Ted. 2026. "Protect the Earth—Build More Data Centers." *The Wall Street Journal*, September 1, 2026. https://www.wsj.com/opinion/protect-the-earthbuild-more-data-centers-5daf794a.

Rust, Max, and Will Parker. 2026. "These Cities and States Are Taking Aim at Data Centers." *The Wall Street Journal*, April 7, 2026. https://www.wsj.com/tech/ai/these-cities-and-states-are-taking-aim-at-data-centers-3b98adf1.

Schwartzel, Erich. 2026. "The 'Country Hicks' Who Refused $26 Million from an AI Data Center." *The Wall Street Journal*, August 15, 2026. https://www.wsj.com/tech/ai/ai-data-center-rural-america-backlash-c0af4e16.

Uberti, David. 2025. “What Happened When Small-Town America Became Data Center, U.S.A.” *The Wall Streat Journal*, November 3, 2025. https://www.wsj.com/us-news/what-happened-when-small-town-america-became-data-center-u-s-a-410f25e9.

UNU-INWEH (United Nations University Institute for Water, Environment and Health). 2026. *Environmental Cost of AI’s Energy Use: Carbon, Water and Land Footprints*. https://doi.org/10.53328/inr26rma002.

Yale School of the Environment. 2026. *Data Centers, Urban Heat, and AI Growth*. https://environment.yale.edu/news/article/data-centers-urban-heat-and-ai-growth.

Casualty Actuarial Society

4250 North Fairfax Drive, Suite 650
Arlington, Virginia 22203

casact.org