Closing the Gap.

What AI readiness actually takes in reinsurance

Toucanberry Tech, September 2026.

TOUCANBERRY.COM

Contents.

Executive summary 03

The problem 05

Why this matters more now 07

Figure one: one number, end to end 09

Step one: bring the data together 10

Step two: make it trustworthy 11

Step three: put it into the workflow 12

Five questions to ask about one workflow 15

What closing the gap makes possible 16

Sources and notes 17

Toucanberry Tech, September 2026.

Executive summary.

Whether AI works inside a reinsurer is decided long before anyone chooses a model. The gap is between AI ambition and the operational foundations needed to support it. This report sets out what closing it involves, and how to work out what it would take in your own operation.

AI is on both sides of that gap - it is what most firms want to deploy next, and what fails fastest when the data underneath cannot be relied on. Three steps close it, each one making the next possible and putting a new capability within reach, ending with AI running inside the tools the team already uses.

The same work pays twice - it shortens the close cycle and takes manual work out of reporting as soon as it is done. It is also the foundation that AI in any core reinsurance workflow depends on.

Our January report, *AI Readiness in Bermuda Life Reinsurance*, called that foundation operational readiness. We found 99% of insurers had AI initiatives underway, and far fewer with the operational foundations to support them.

THE EVIDENCE SINCE JANUARY

Evidence published since then points the same way. Capgemini found 42% of P&C insurers had set no AI KPIs at all, and 60% were still in exploration or proof of concept. An AM Best survey of more than 150 rated carriers identified data readiness as one of the largest barriers to implementation. EIOPA, the EU's insurance regulator, found around two-thirds of European insurers using generative AI, with most of it still at pilot stage.

That foundations matter is no longer the argument. The question is how to move from tests and pilots to operational AI at scale.

WHY NOW

Two things make closing the gap more urgent. Agentic AI has moved from demonstrations into serious deployment conversations, and it depends on the foundation most reinsurers have still not built.

And regulators have begun setting out what they expect, most recently the Bermuda Monetary Authority, whose August consultation on the responsible use of AI wants firms able to reconstruct how a material decision was reached. That is difficult to evidence when the controls themselves are spreadsheets.

So what has to be true before AI can work reliably inside a reinsurer? Three things, and the order matters.

01 Bring the data together.

Data arrives from cedants, asset managers and investment houses in every format there is, and most of it moves by hand in Excel. Until it sits in one governed set, nothing can be asked of it reliably.

Once it does: someone can ask a question of the data in plain English and get an answer back.

02 Make it trustworthy.

Validation, reconciliation and control checks run automatically across the whole lifecycle of the data, including after the actuarial modelling. Numbers validated before anyone acts on them, exceptions visible, and an audit trail to prove it.

Once that's done: bad data gets caught before it reaches a report.

03 Put it into the workflow.

Regulatory and financial reporting produced automatically. Exceptions routed, owned and closed out with a record.

Once that's in place: AI starts doing real work inside the tools the team already uses.

The result is a reinsurer that can grow the book without growing the operation to match.

CASE STUDY

Quarter-end close: 15+ days to 3 days.

The close is where every problem in this report converges: fragmented feeds, controls on spreadsheets, reporting assembled by hand. One Bermuda reinsurer now runs its core quarter-end close in three working days.

Sophisticated modelling, manual everything else.

Our January report set out what that looks like in practice, and none of it has moved since. Firms model complex longevity scenarios while coordinating daily work through spreadsheets and email. They price new structures while data is still moved manually between systems. They stress-test portfolios under multiple scenarios while reconciling the exact numbers by hand across three platforms.

The question is what that gap looks like inside a reinsurance operation, because that determines what has to be built.

More than one version of the same number

Almost every reinsurer we work with has the same problem in some form.

Data arrives from multiple sources in spreadsheets. It gets combined, transformed and stored in more than one place, so more than one version of a figure exists. The control checks that establish which version to trust run on spreadsheets too, with formulas nobody wants to unpick. Then the reporting gets assembled on top of all that, largely by hand.

The data still gets used. Actuaries run their experience studies; reserving teams work the triangles. But each one starts from scratch - pulling records from system after system, cleaning them by hand, reconciling definitions that never match - before the analysis can begin.

What's missing is a repeatable way to establish which version of a figure is authoritative.

What it costs.

Everything downstream inherits the version problem. Regulatory and financial reporting sits on numbers that have been checked in fragments rather than through one process end to end. A figure gets reported, accepted and booked, and a later control cycle establishes that a different version was the right one. Someone goes back and revises a number that had already been treated as settled.

One firm put a number on the cost with a time audit of its actuarial team. Leadership had known something was wrong. What they had not known was that it came to two days a week, per actuary, spent moving data rather than analysing it.

Projects are the other cost. In our own discovery work, planned builds have stopped or changed direction at discovery, once the data underneath was examined.

A proposed valuation dashboard was replaced by import error detection, because the errors were upstream of anything a dashboard could show. Front-end improvements were set aside until the flow of data beneath them was fixed. A treaty data intake process had to be cleaned up before anything else could start.

TIME AUDIT

~40% of one actuarial team's hours went on manual data work.

Retrieving, reformatting, reconciling and preparing model inputs, at a single firm.

Why this matters more now.

Two things have changed since the January report.

01 Agentic AI

Agentic AI has moved from demonstrations into serious deployment conversations. But an agent inherits the version problem. Point it at disconnected systems and it hits the same wall, then carries on. Experienced staff spot inconsistencies that an agent, without the right controls, will simply process. People learn where the data cannot be trusted and route around it; an agent has no such instinct.

It needs the data joined up, validated at the point of entry and traceable as it moves, or it automates the mistakes faster than anyone can catch them. And you cannot fix it after the fact. An agent that calculates across data nobody has governed produces a figure shaped by whatever it was given, and no amount of checking the output afterwards undoes the calculation.

40%+ of agentic AI projects forecast to be scrapped by the end of 2027.

Undone by rising costs, unclear value and controls that were never really there.

GARTNER, 25 JUNE 2025

02 The regulators

Regulators have been moving in the same direction. The IAIS, the global standard-setting body for insurance regulators, published its application paper on supervising AI in 2025. EIOPA issued an opinion on AI governance and risk management. The NAIC, which represents US state regulators, has a model bulletin on insurers' use of AI now adopted in 25 states.

Most directly, the BMA published its consultation on the responsible use of AI on 14 August, with comments closing 30 October. It creates no new obligations or approval regime, instead describing how existing expectations apply.

Three of those expectations bear directly on the work described in the steps that follow.

01	For material or higher-risk use cases, records sufficient to reconstruct a decision: inputs, outputs, human review, overrides and the basis for the final call. In practice, traceability and data lineage.
02	Shadow AI as a governance challenge, including AI arriving embedded in third-party software without anyone deciding to adopt it.
03	Managing both the risks of using AI and, where relevant, the operational risks of failing to modernise.

In practice that means governing access to inputs and outputs from the outset rather than reconstructing it later, and logging which outputs were served, so a decision can be traced back to what the model actually produced when it was taken.

AG55, the NAIC's Actuarial Guideline 55, requires US cedants to test the adequacy of reserves they cede to reinsurers. It has already been through its first reporting cycle.

The NAIC's early review flagged recurring difficulties: explaining reserve movements, substantiating reserves held by the reinsurer, documenting assumptions and testing sensitivities. Each depends on being able to trace where a number came from and stand behind the assumptions underneath it.

1 April 2026

First AG55 filings due, around 80 ceding companies and 100 treaties

14 August 2026

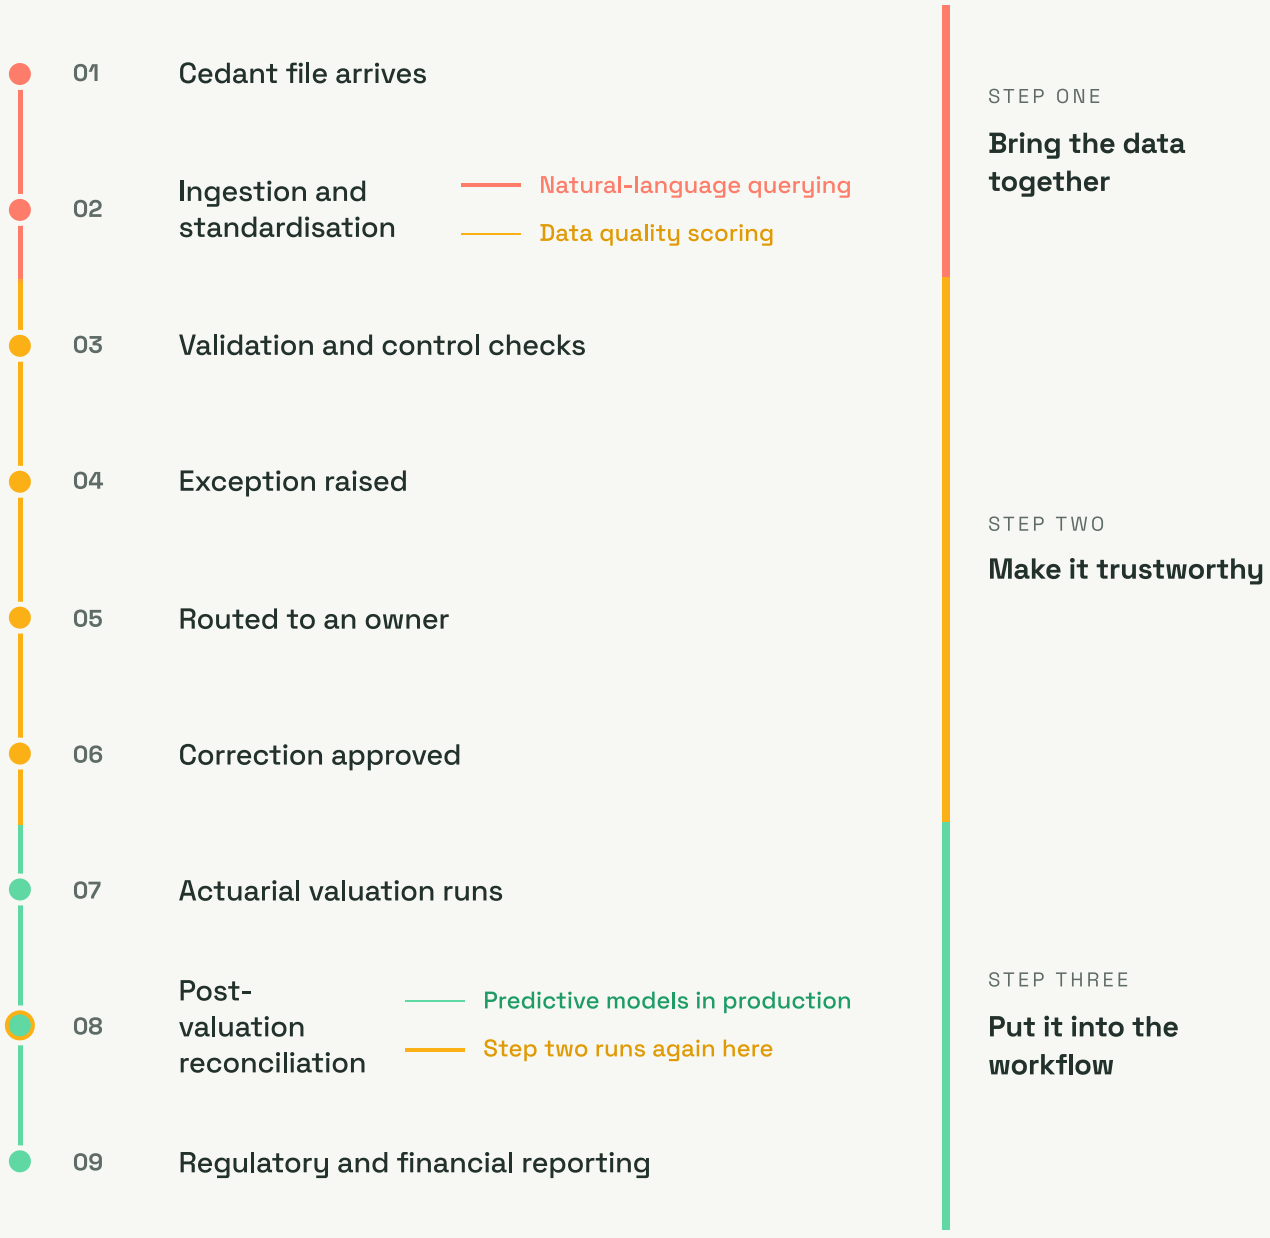
BMA consultation published

30 October 2026

Comments close

One number, end to end.

This is what happens to a cedant file between arriving and appearing in a report. Nine stages, most of them handled by hand today. The three steps in this report each cover a section of it, and each one makes a new capability possible.



Validation and control run across the whole lifecycle of the data, not only at ingestion.

01 Bring the data together.

WHERE THIS SITS IN THE FLOW

FIGURE ONE, STAGES 01-02



Cedant file arrives - ingestion and standardisation

The first job in getting AI-ready is to bring the data into one governed set, from every source, in a form the whole business can use.

That means ingesting whatever arrives, tracking what is expected, what has been received and what is overdue, and turning it into one set of numbers. In practice that resolves into a handful of structured views: liability data, asset data, holdings, investment data. The systems you already run stay where they are - the actuarial environment, the asset tools, the regulatory reporting framework. What changes is that they are connected.

None of this has to be the whole enterprise at once. The same sequence works inside a single workflow or function, which is usually where it starts.

The hardest part is usually the incoming files - seriatim data in life, loss and exposure files in P&C - arriving from cedants in whatever format each one uses. Turning that into clean, validated, actuarial-ready data has traditionally been manual work. At several reinsurers we work with, that now happens automatically.

It's also the step that puts the AI capabilities everyone wants within reach.

WHY STEP ONE ISN'T ENOUGH

Getting it into one governed set is not the same as being able to rely on it. In one engagement, a client had been processing files for months before anyone spotted that a change to the incoming format had corrupted downstream reports. The numbers simply stopped being right, but no part of the process was designed to notice.

WHAT THIS PUTS WITHIN REACH

Natural-language querying.

Someone asks a question of reinsurance data in plain English and gets an answer back, without commissioning an extract or joining a reporting queue. It needs one governed set underneath it, consistent enough to query safely.

02 Make it trustworthy.

WHERE THIS SITS IN THE FLOW

FIGURE ONE, STAGES 03-06, AND 08



Validation and control checks - exception raised - routed to an owner - correction approved. Stage 07 is the valuation itself; step two resumes afterwards.

Consolidation on its own does not make data reliable.

In our client conversations, the concerns that recur are accuracy, compliance and penalties. Get those wrong and every new book adds risk as well as revenue. Step two also decides whether an agent can be trusted, because an agent acts on whatever it is given.

The work is validation, reconciliation and control checks running automatically across the whole lifecycle of the data, including after the actuarial modelling rather than only at ingestion. Numbers get checked before anyone acts on them, exceptions surface rather than staying buried, and there is an audit trail to prove it.

At one reinsurer, validated data now lands in staging built for actuarial logic, with QA checks and audit trails, before the reserving model runs at all. Reserving cycle time fell by more than half.

Done properly, that changes what the cycle is for. Routine checks run automatically, leaving genuine exceptions - a source-system error, a late adjustment, a point of accounting judgement - to be routed to an owner, investigated and closed out with a record.

WHAT THIS PUTS WITHIN REACH

Data quality scoring.

Every incoming feed profiled against what it should look like, with drift flagged before it reaches a report. It needs historical baselines for each feed, and those only exist once the data sits in one governed set.

03 Put it into the workflow.

WHERE THIS SITS IN THE FLOW

FIGURE ONE, STAGES 07-09



Actuarial valuation runs - post-valuation reconciliation - regulatory and financial reporting

At this point the operational payoff becomes visible, and AI can start doing useful work on the foundation that pricing and deal evaluation will eventually run on too. Reporting stops being assembled by hand, people get answers from the data without waiting on a reporting queue, and models run against live data inside the tools the team already uses. All of it rests on the two steps before it.

In practice three things change. A correction is made once, at source, and flows through to every report that uses it, instead of being patched into each spreadsheet. A question gets answered from the governed set rather than from a fresh extract. And an exception has an owner, a status and a record, so the quarter-end close becomes a list of things to decide rather than things to find.

Step three also forces a decision firms often leave open. Work at this stage changes how reporting gets produced and who answers for it, so it needs a senior business owner, often the COO or CFO. Where these are run purely as technology projects, the technology tends to get delivered and the workflow around it does not change.

WHAT THIS PUTS WITHIN REACH

AI in the tools people already use.

Predictive models running alongside the governed data rather than in a separate reporting layer, monitoring lapse behaviour against its drivers and flagging when actual experience begins to diverge from assumptions. It needs both earlier steps: data that is joined up, and controls that establish which version of it is right. Ours run live on governed data, with the validation, explainability and audit trail a regulated balance sheet needs.

Three examples of what changes.

EXAMPLE ONE

Reporting stops being assembled.

At one Bermuda reinsurer, the core quarter-end close now runs in three working days, down from more than fifteen. The teams still review, still exercise judgement, still sign off. They just spend the week reading the numbers rather than assembling them.

EXAMPLE TWO

Analysis becomes something people do for themselves.

Actuaries who had never written a line of code are generating working Python against their own production data. That becomes possible because the underlying data is consistent enough to be queried safely. Getting an answer no longer means waiting for a specialist team.

EXAMPLE THREE

Models get built where the work already happens.

On one engagement, years of lapse experience had never been pulled into one place. Once it was cleaned, standardised and connected to the tools the actuaries already used, we built predictive models that watch lapse behaviour against drivers such as Treasury yields and flag the moment experience starts to drift from assumption, running alongside the governed data rather than in a separate reporting layer. Nothing about that required new data. The firm had owned every record for years.

Growth without the operation growing.

The work does not have to be taken as one programme. Firms can start with the workflow or capability that solves the problem in front of them. From there the operating model can vary too: some capabilities run as a managed service, others go directly into teams' hands, and a firm can have a mixture. The same foundation supports any of them.

None of this replaces the question of whether the process is the right one. Automating a workflow nobody has examined gives you the same workflow, faster. The foundation is what makes that examination possible, because you cannot redesign what you cannot see.

One caution on released capacity, which tends to be oversold. Time savings are real, but what happens to the time is a separate question. Anyone who lived through IFRS 17 or LDTI, the accounting changes that forced multi-year system programmes across the industry, knows what happened to the capacity those systems were meant to release. Most of it was absorbed by running the new reporting. Removing manual work produces available hours and a decision about them. Decide in advance where released capacity should go, whether that is supporting additional books, deeper analysis, or work that is currently deferred.

CASE STUDY

One book to seven, on a team of around fifty.

One Bermuda reinsurer went from one book of business to seven in under four years, while also shortening the close. Clean data moved between systems without anyone carrying it, so manual work did not grow with the book.

BERMUDA LIFE REINSURER · TOUCANBERRY ENGAGEMENT

Five questions to ask about one workflow.

You do not need to assess the whole operation at once. Take one workflow, the one you most want AI working in, and answer these five questions. The answers show whether it is ready and, if not, what needs fixing first.

01 AI

If you pointed an agent at this workflow tomorrow, what could it act on safely, and what would it get wrong?

The next four tell you the answer.

02 DATA

Can anyone say which version of a figure is the authoritative one, and what does it take to establish that?

03 CONTROLS

If an incoming feed changed shape without warning, how long before anyone noticed?

04 WORKFLOW

How many manual handoffs sit between data arriving and the output of this workflow?

05 GOVERNANCE

If a regulator or an auditor challenged a material figure, could you reconstruct how it was reached, end to end?

What closing the gap makes possible.

Whether AI works inside a reinsurer is decided long before anyone chooses a model. The work starts with the operation rather than the technology - a reinsurer with joined-up data, automated controls and repeatable workflow closes faster, spends less time assembling and reconciling information, and gives actuaries direct access to data they can trust. Those improvements arrive first, whether or not a single AI model has been chosen. They are also the foundation AI depends on.

Technology is what makes that operating model repeatable at scale: the ingestion, the controls, the models that run in production. It also amplifies whatever it is applied to. Put AI on top of a governed operation and it extends what that operation can do. Put it on top of a fragmented one and it makes the same mistakes, faster and more often.

That is why the sequence matters more than the tooling. Natural-language interfaces, predictive models and agents become materially more valuable running on governed data, inside workflows with clear ownership and an audit trail. And the operation underneath them is the one cedants already score, on execution and support, in NMG's Business Capability Index. Taken together, that is a reinsurer that can grow the book without growing the operation to match, which is the only way to chase growth and hold margin at the same time. The same foundation is what pricing and deal evaluation run on too, once the operation can carry the business it wins.

None of it requires an enterprise programme to begin. Take one important workflow, apply the five questions, establish where the constraint sits, then work the sequence. Bring the data together, make it trustworthy, and put it into the workflow.

Sources and notes.

EXTERNAL

Capgemini, *World Property & Casualty Insurance Report 2026*, 5 May 2026. Survey of 344 senior executives, 809 employees and 1,113 policyholders.

AM Best, *Artificial Intelligence Appears to Be Ready, But Most Insurers Are Not*, 27 April 2026. Survey of more than 150 rated carriers and assessed MGAs.

EIOPA, report on generative AI in European insurance, February 2026. 347 undertakings across 25 countries.

IAIS, *Application Paper on the supervision of artificial intelligence*, July 2025.

EIOPA, *Opinion on artificial intelligence governance and risk management*, August 2025.

NAIC, *Model Bulletin on the Use of Artificial Intelligence Systems by Insurers*, adopted by 25 states as at July 2026.

NAIC, Actuarial Guideline 55 (AG55) review materials, June 2026. First filings due 1 April 2026, covering approximately 80 ceding companies and 100 treaties.

Gartner, *Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027*, press release, 25 June 2025.

NMG Consulting, *Global Life & Health Reinsurance Study*, Business Capability Index. Annual survey of approximately 1,500 insurance executives across more than 50 countries, measuring insurers' assessment of reinsurer capability, execution and support.

Bermuda Monetary Authority, *The Responsible Use of Artificial Intelligence in Financial Services in Bermuda*, consultation paper, 14 August 2026. Paragraph 111 (records sufficient to reconstruct material decisions); paragraph 32 (shadow AI as a governance challenge); paragraph 2 (AI embedded in third-party software, and the risks of failing to modernise).

TOUCANBERRY CLIENT ENGAGEMENTS

Time audit finding (40% of actuarial hours on manual data tasks): single-firm engagement, January 2026 report.

Quarter-end close reduced from more than fifteen days to three working days: Bermuda life reinsurer engagement.

One book to seven in under four years on a team of around fifty: Bermuda life reinsurer engagement.

Reserving cycle time reduced by more than half: automated pipelines into Axis with staged, validated data and QA checks, reinsurer engagement.

Discovery findings (valuation dashboard, front-end improvements, treaty data intake): Toucanberry discovery engagements.

Lapse behaviour models, seriatim ingestion, data quality scoring and natural-language querying: client engagements and lab work.

Toucanberry Tech works with reinsurers to connect the data, automate the controls and build the workflow around them. If you would like to compare that against one of your own workflows, get in touch at laith@toucanberry.com.

SEPTEMBER 2026

Closing the Gap.

What AI readiness actually takes in reinsurance

Toucanberry Tech, September 2026.