Toucanberry Tech.

AI Readiness in Bermuda Life Reinsurance.

Building Scalable Operations To Support AI

Contents:

Executive Summary

Bermuda's life reinsurers are beginning to invest in AI, but most lack the operational foundation to use it effectively.

The sector manages over $1 trillion in assets, with 80 per cent of assets under management running on collateralised structures. Growth has been extraordinary.

Regulatory pressure is intensifying, with AG55 reports due in April 2026 and BMA enhanced disclosure through 2025 and 2026. Meanwhile, political scrutiny from US and EU regulators continues.

Globally, 99% of insurers report AI initiatives underway, though Bermuda's life reinsurance market remains at a much earlier stage. The constraint isn't ambition or technology availability; it's operational readiness.

99%
Of insurers have AI initiatives underway

VS

<20%
Can complete month-end close in under one week

AI ambition far outpaces operational capability.

This report draws on regulatory analysis, industry benchmark research, conference insights from BILTIR and SIRC 2025, and Toucanberry's operational experience with Bermuda life reinsurers.

Three critical findings:

1. Data infrastructure, not algorithms, determines AI success.

Most firms layer sophisticated tools onto fragmented systems.

Actuaries manually reconcile data between platforms. Settlement processes take days instead of hours. Some reinsurers still spend 2-3 weeks after quarter-end just validating data before analysis can begin. Month-end close stretches past two weeks, whilst top performers in adjacent industries finish in five days.

The gap between analytical capability and operational infrastructure is where AI implementations fail.

This doesn't mean waiting until everything is perfect. Three paths move forward in parallel.

- Systematic data infrastructure - building foundations that enable AI at scale.
- Personal productivity AI - deploying generative AI for email, documentation, and analysis with clear governance frameworks.
- Structured pilots - experimenting with AI on specific processes where workflows are mapped and data quality is understood.

2. Regulation is forcing overdue modernisation.

AG55 requires cash-flow testing and attribution analysis that most firms can't currently produce.

BMA disclosure requirements demand real-time visibility into asset mix and ALM strategy.

These aren't concerns about a distant future - they're impending 2026 deadlines. Against that backdrop, firms treating compliance as separate from operational transformation will find themselves building the same capabilities twice. Those who recognise regulation as a catalyst for necessary infrastructure investment gain a competitive advantage whilst meeting requirements.

3. Collateralised structures create unique opportunities.

80 per cent of Bermuda's life reinsurance uses collateralised structures in which assets are held to secure obligations - either funds-withheld (cedant holds the assets) or modified coinsurance arrangements.

Cedants increasingly evaluate partners based on reporting speed, data quality, and real-time transparency.

Capital strength alone no longer differentiates. Every major player has robust balance sheets and sophisticated underwriting. Rating agencies are increasingly focusing on operational resilience - the ability to scale without proportional headcount growth, to close books quickly, and to provide real-time transparency.

The operational imperative

Leadership needs to understand whether current operations can actually support the technology they're buying. Most firms only discover this when implementation starts: the systems can't talk to each other, data doesn't flow cleanly, and the pilots that worked in isolation fail in production.

Firms that can provide real-time collateral reporting and stress testing won't just comply better - they'll win more business.

The typical pattern amongst the firms pulling ahead is that their starting point wasn't to ask "which AI tools?". They started with operations. They consolidated data infrastructure. They automated workflows. They established governance frameworks with clear ownership. They measured what matters - cycle times, manual touchpoints, exception rates.

In session polling at BILTIR (Bermuda International Long-Term Insurers and Reinsurers) 2025, operational leaders ranked data quality as the number-one barrier to AI adoption, ahead of technology costs, vendor selection, and regulatory uncertainty.

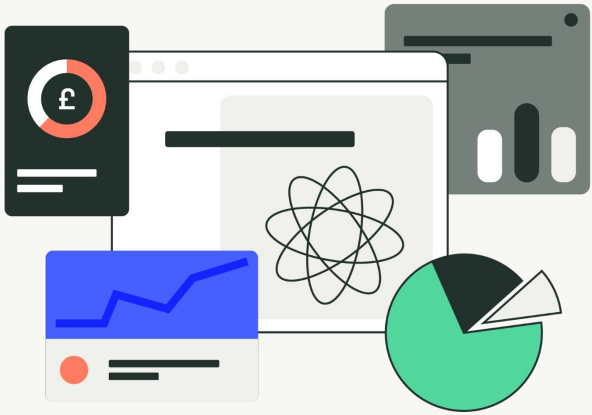

1. What AI Readiness Actually Means

Type	What It Does	Reinsurance Application
Machine Learning	Pattern recognition from historical data	Risk scoring, mortality studies, pricing optimisation
Process Automation	Rule-based task execution	Data validation, reconciliation, straight-through processing
Generative AI (LLMs)	Text generation and document analysis	Regulatory filing extraction, treaty review, email drafting
Agentic AI	Multi-step autonomous workflows	Early exploration stage - not yet in production use

There's a series of unusual disconnects at play in the industry.

- Firms can model complex longevity scenarios whilst coordinating daily work through spreadsheets and email.
- They price exotic structures whilst waiting days for data to move between systems.
- They stress-test portfolios under multiple scenarios whilst manually reconciling the exact numbers across three platforms.

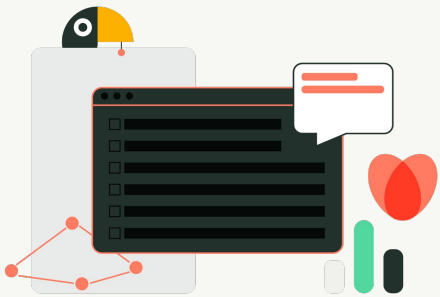

Three dimensions determine whether AI contributes to growth or remains a pilot project.

1. Workflow readiness

Identify where manual coordination breaks workflows.

- Data extraction from cedant files.
- Validation against expected formats.
- Reconciliation between investment and actuarial systems.
- Month-end close processes that should take days but stretch to weeks.

Manual coordination isn't a technology gap - it's a workflow design gap.

Firms running efficient operations have systematically mapped their processes. They've identified handoffs between functions. They've rebuilt coordination as system-driven rather than people-dependent.

Process mapping means understanding work at the task level.

- Which tasks are routine data validation versus judgment calls?
- Which require human expertise versus following rules?
- Which create bottlenecks because information moves through email instead of systems?

Identifying tasks that AI can support starts with understanding which tasks exist in the first place.

The best reinsurers redesigned processes from the start for automation rather than retrofitting legacy practices. Gen Re, for example, achieved 69 per cent straight-through processing for underwriting applications in group medical lines through workflow redesign, not better AI.

Industry benchmark studies show that top-quartile performers complete month-end close in 4.8 days, compared with 10+ days for the bottom quartile. Again, that's a workflow design benefit, not a reflection on their tech budget.

2. Data architecture

Legacy actuarial systems weren't built for real-time operations. They expect quarterly data loads, batch processing, and manual validation loops. Adding AI tools on top doesn't fix the underlying constraint.

Modern platforms provide continuous data pipelines. Investment positions update automatically. Cedant submissions flow through structured validation. Actuarial models run on current data, not last quarter's snapshot.

The shift from quarterly to continuous is, therefore, an operational necessity. Firms operating with 90-day visibility cycles can't compete with those operating with daily visibility. These questions reveal whether your data architecture can support AI, regardless of which tools you buy:

1. How long does your month-end close actually take?
2. How many manual touchpoints are there between raw data and regulatory filings?
3. Where do version control issues still happen?

3. Governance frameworks

The Bermuda Monetary Authority (BMA) released its AI governance discussion paper in 2025. The approach is principles-based, proportionate, and substantive. Model validation requirements under the Scenario-Based Approach demand independent review, documented assumptions, and systematic oversight.

Board-level ownership matters here. The firms succeeding with operational transformation have clear C-suite accountability.

Research across financial services firms shows that those with centralised AI governance achieve 70 per cent production success rates, whilst federated models achieve 30 per cent. Centralised governance means unified oversight under a single executive structure - typically led by the COO or CFO - rather than fragmented departmental pilots.

IT-only initiatives fail because they can't address workflow and governance challenges that span the organisation.

75 per cent of banks now have formal governance committees. 79 per cent say they'd prioritise governance if they were to restart implementation.

Organisational design matters more than technology spending here.

Teams need the authority to make decisions without requiring consensus from every affected department. They need to measure cycle times, exception rates, and manual touchpoints - not just count AI pilots launched.

Workforce AI adoption

Infrastructure readiness only tells half the story. How people actually work with AI tools matters as much as whether the systems can integrate.

Analysts using generative AI for research and documentation report 60-70 per cent time savings on routine tasks. Already, actuaries can use large language models to parse complex regulatory text and extract requirements. Investment teams deploy AI to screen deal flow and flag portfolio risks faster than manual review allows.

But this organisational AI use creates its own readiness questions. Two patterns are emerging.

First, shadow AI - employees using personal ChatGPT accounts, Gemini, or other tools without organisational oversight. This creates information security risks and inconsistent quality, but it's inevitable and already happening across the sector.

Second, systematic workforce AI adoption - organisations providing approved tools, establishing validation protocols, and training teams on appropriate use cases. Email drafting, meeting summarisation, regulatory text analysis, and document review don't require rebuilt data infrastructure. They require clear guidelines and governance.

This creates practical questions:

? *Do teams have clear guidelines on when to use AI versus when human judgment is required?*

? *Are outputs being validated systematically?*

? *Is there governance around what data can be fed into external AI services?*

The firms moving fastest deploy personal productivity AI with governance whilst building systematic data infrastructure in parallel. Both matter. Both move forward together.

CASE STUDY: Building AI-ready ops

A Bermuda-based life reinsurer, part of a global insurance group

This reinsurer demonstrates this approach across both infrastructure and organisational AI use. They experienced problems common to those in the Bermuda ecosystem.

- Rich data trapped in disconnected systems.
- Month-end close taking 15+ days.
- Manual processes blocking any automation attempt.

Rather than prioritising the procurement of new tools, they started with workflow redesign.

- They decoupled data pipelines from legacy actuarial systems, creating a unified data architecture storing over **2.2 million** data points.
- Settlement statement processing that once took days **now completes in minutes.**
- They automated **80 per cent** of core business processes before introducing any "AI."

Machine learning models were deployed incrementally - integrated where clean data and systematic processes made them useful.

The result: when compliance changes e.g. AG55 requirements arrive, they have the operational foundation to comply efficiently.

More importantly, they can provide cedants with real-time portfolio updates - a competitive differentiator in their segment. They've grown from 1 to 7 books in under four years with approximately 50 people. Traditional scaling would have required a headcount three times that.

AI readiness came from operational discipline, not vendor selection.

The Operational Foundation in Numbers

1 → 7	books under management in under 4 years
~50	people managing 7 books (traditional model would require 150+)
15+ days → 3 days	month-end close cycle time (80% reduction)
80%	of core business processes automated before any "AI"
Weeks, not months	necessary to achieve compliance readiness

AI readiness came from operational discipline, not vendor selection.

The pattern repeats across firms succeeding with AI.

- They measured the current state first - how long does month-end close actually take?
- How many manual touchpoints exist in critical processes?
- Where do coordination failures happen?

Then they fixed the operations. Consolidated data. Automated workflows. Established governance. Built measurement into the process.

As operations improve, AI becomes progressively more useful - deployed incrementally where clean data and systematic workflows can support it.

AI Readiness Assessment: Where Does Your Firm Stand?

Even if current operations still rely on manual coordination, new processes can be designed with AI-readiness in mind - embedding automation hooks, data validation, and governance from day one.

Dimension	Level 1: Not Ready	Level 2: Partially Ready	Level 3: AI-Ready
Workflow Readiness	• Manual coordination • 10+ day close • Email-driven processes	• Some automation • 5-10 day close • Hybrid manual/system	• System-driven • < 5-day close • Exception-based management
Data Architecture	• Quarterly loads • Batch processing • Manual reconciliation	• Periodic updates • Some real-time access • Partial validation	• Continuous pipelines • Real-time visibility • Automated validation
Governance	• IT-led initiatives • No clear ownership • Pilot stage only	• Nominated leaders • Some coordination • Selected production	• C-suite ownership • Clear authority • Systematic measurement

2. Common Failures

MIT research shows 95 per cent of GenAI pilots fail - not from technical limitations, but because organisations avoid the operational friction required to make them work.

Five patterns explain this trend across the reinsurance industry. These aren't technology problems - they're organisational design failures.

Failure Pattern	What It Looks Like	Result
Pilots without workflows	AI model predicts settlement timing, but process uses email coordination between three teams, manual data entry, phone calls	Expensive platforms sitting unused in production. Technology works, operations don't support it
No operational measurement	Teams know it's "too long" or "too manual" but have no baseline metrics	Improvement becomes subjective - teams argue with zero data to support claims
Wrong ownership structure	IT leads because it's technology. Business doesn't own implementation. No one owns integrated outcome	Model works technically but never reaches production after six months
No documented strategy	Different teams pursue different tools - actuarial, risk, operations each working independently	Each pilot becomes maintenance burden. Fragmented capabilities, no common standards
Build vs buy as procurement	Platform selection treated as vendor comparison by IT procurement team	Miss strategic implications of architecture choices

The common thread: firms tried to buy their way to automation without building the systematic workflows that make automation possible.

Common Failures in AI Adoption

MIT research shows 95 per cent of GenAI pilots fail - not from technical limitations, but because organisations avoid the operational friction required to make them work.

Five patterns explain this trend across the reinsurance industry. These aren't technology problems - they're organisational design failures.

Failure Pattern	What It Looks Like	Result
Pilots without workflows	AI model predicts settlement timing, but process uses email coordination between three teams, manual data entry, phone calls	Expensive platforms sitting unused in production. Technology works, operations don't support it
No operational measurement	Teams know it's "too long" or "too manual" but have no baseline metrics	Improvement becomes subjective - teams argue with zero data to support claims
Wrong ownership structure	IT leads because it's technology. Business doesn't own implementation. No one owns integrated outcome	Model works technically but never reaches production after six months
No documented strategy	Different teams pursue different tools - actuarial, risk, operations each working independently	Each pilot becomes maintenance burden. Fragmented capabilities, no common standards
Build vs buy as procurement	Platform selection treated as vendor comparison by IT procurement team	Miss strategic implications of architecture choices

The common thread: firms tried to buy their way to automation without building the systematic workflows that make automation possible.

3. Where AI Delivers Competitive Advantage

Most reinsurers built operations for periodic reviews, but now have to deliver continuous oversight using the same manual processes.

You can see this trend developing in other markets as well.

Zurich Insurance has deployed an agentic AI platform in Commercial Insurance that compresses document handling from hours to minutes and is targeting a step-change in quote throughput and renewals productivity. It works because the data and workflow are systematic - underwriters get real-time support inside the process, not in a sidecar tool.

Aon's Impact Forecasting group uses AI to generate stochastic event sets for tropical cyclones and to drive automated event-day loss estimation. The outputs flow straight into pricing, exposure management, and claims prioritisation because the operational plumbing - data standards, live ingestion, predefined decision points - was already in place.

Swiss Re deployed ClaimsGenAI for automated claims analysis, identifying millions in potential fraud savings in the first year. These applications succeed because the operational foundations - data standards and systematic workflows - were already in place.

The lesson is the same: when data is structured and workflows are systematic, AI delivers measurable results. These projects succeed not because of superior algorithms, but because the operational foundations were in place.

Life reinsurance isn't catastrophe modelling, but the principle is identical - when data lineage and workflows are defined, AI delivers usable outputs at business speed.

- **Real-time asset-liability matching as market conditions change.**
- **Automated collateral calculations update as positions move.**
- **Predictive stress testing before regulators require it.**
- **Attribution analysis decomposes reserve changes across asset classes, rates, and assumptions.**

These applications solve defined problems where manual processes demonstrably can't keep pace with reporting demands. You're not exploring what AI might do - you're solving specific problems where current approaches fail.

Cedants are changing evaluation criteria. Some now ask during due diligence: how often can you provide portfolio updates? What's your cycle time for stress testing? How do you handle quarterly attribution analysis?

Not compliance attestations. Actual operational capabilities.

Firms that can demonstrate working systems are winning business. Better operations reduce cedant uncertainty, which builds trust. Trust leads to more business. More business justifies operational investment. The cycle compounds.

Operational capability is becoming a selection criterion, not just pricing.

> *You're not exploring what AI might do - you're solving specific problems where current approaches fail.*

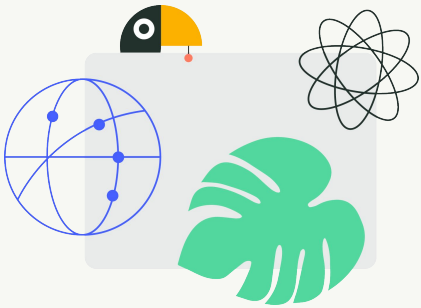

4. AG55 as Catalyst

AG55 is one regulatory catalyst among broader structural pressures - including the BMA's AI governance framework and enhanced disclosure requirements - but it creates the immediate deadline for Bermuda life reinsurers.

These aren't obstacles. They're forcing functions for operational investment firms should have made anyway. AG55 represents the first major US regulatory requirement specifically targeting asset-intensive reinsurance to offshore jurisdictions.

First reports are due April 1, 2026. The requirement is initially disclosure-only, but it creates operational demands that most reinsurers haven't addressed.

Asset-level look-through when portfolio visibility is limited

Attribution analysis explaining reserve decreases.

Cash-flow testing using multiple scenarios.

Documentation of assumptions.

The technical requirements are achievable. The operational challenge is doing this systematically, repeatedly, and efficiently.

Firms that still close their books 2-3 weeks after quarter-end will find AG55 compliance expensive and error-prone. Every manual reconciliation step introduces delay and risk. Every spreadsheet-based validation creates key person dependency.

The attribution analysis requirement is particularly revealing. You need to trace exactly how reserves change due to differences in mortality assumptions, interest rates, lapses, expenses, and asset performance. Then document it. Then defend it to regulators who will question the methodology.

Doing this manually for one treaty is possible. Doing it systematically across an entire portfolio whilst maintaining audit trails and version control requires operational infrastructure that most firms don't have.

BMA's enhanced disclosure requirements follow similar logic.

- More granular asset holdings.
- Product-level reserves and durations.
- ALM strategy and liquidity stress testing results.
- Testing expected through 2026.

The expectation is real-time operational capability, not point-in-time compliance efforts.

Firms investing now in data architecture and workflow automation will find compliance straightforward. Those waiting will discover their competitors can report faster, with higher confidence, at lower cost.

Bermuda's regulatory approach - business-friendly yet substantive - creates room for genuine operational improvement rather than just checking compliance boxes.

The Implementation Reality: Regulatory Deadlines vs Operational Timelines

The regulatory and operational timelines are converging. Firms beginning infrastructure work today will barely meet AG55 requirements; those waiting will miss them.

REGULATORY MILESTONES	Q4 2025 AG55 effective date	April 1, 2026 ⚠️ First reports due	Through 2026 Partial manual compliance	2027 Full automation expected
OPERATIONAL RESPONSE	Q1 2026 Assessment & Strategy	Q2 2026 Design & Procurement	Q3-Q4 2026 Build & Test	2027 Ready for automated cycle

⚠️ ***Platform implementations take 6-18 months. Starting Q1 2026 means being ready for 2027, not 2026.***

> *Firms investing now in data architecture and workflow automation will find compliance straightforward.*

5. What Leadership Should Do

Five questions for your next board meeting:

→ Can we actually see our data in real-time, or are we waiting for quarterly reports?

If quarterly, you're not AI-ready, regardless of the tools you purchase. Real-time visibility is table stakes for modern operations. If actuaries need to wait for quarter-end to run scenarios, if risk teams can't access current positions without requesting data extracts, or if cedants ask for portfolio updates and the answer is "we'll have that for you next week," you're losing ground to the competition.

→ How many manual touchpoints exist between data ingestion and regulatory reporting?

If the answer is "don't know," start there. Every manual step introduces delay, error risk, and scaling constraint. Map the actual process - not the intended process, the real one. Most firms discover they have 50+ manual touchpoints in processes they thought were largely automated.

→ Who owns operational transformation - IT, risk, operations, or CFO?

If the answer is unclear, the transformation won't work. This requires business leadership, not just technology implementation. If it's IT alone, they don't own business outcomes. If it's federated across departments, nobody has the authority to change how processes integrate. If it's a steering committee without an executive sponsor, the committee won't make decisions under pressure.

→ What's our month-end close cycle time, and why isn't it faster?

This single metric reveals operational maturity. Top-quartile performers complete month-end in under 5 days, versus 10+ days for bottom-quartile performers. If you're at 10+, you have systematic operational dysfunction blocking automation. Fast close cycles require systematic data flows, automated validation, exception-based management, and processes that don't depend on individual expertise.

→ How will we comply with AG55 attribution analysis requirements by April 2026?

If the answer is "we'll figure it out," you're already behind. The firms ready for AG55 built systematic capabilities over 12-18 months. Can you systematically decompose reserve changes across mortality assumptions, interest rates, lapses, expenses, and asset performance? Can you document the methodology? Can you produce audit trails? Can you do this for your entire portfolio, not just one treaty? Board review of AG55 readiness should happen before Q2 2026 budgets lock.

6. The Bottom Line

Bermuda life reinsurance figured out asset-intensive structures. It can determine the operational infrastructure needed to support them.

But this $1 trillion sector requires capabilities that most firms haven't yet built.

- **Real-time portfolio visibility.**
- **Systematic attribution analysis.**
- **Stress testing that delivers results in days, not weeks.**
- **Reporting infrastructure that supports continuous oversight rather than periodic reviews.**

Better data enables faster decisions. Faster decisions enable more competitive pricing. More competitive pricing wins more business. More business creates more data to learn from.

The technology exists. The constraint is implementation.

Banking and asset management faced similar challenges a decade ago. Firms that rebuilt data infrastructure and systematic workflows before layering AI achieved measurable results - faster close cycles, lower compliance costs, better risk oversight. The playbook exists for reinsurers willing to prioritise foundations over features.

AG55 requirements due in April 2026, and BMA enhanced disclosure through 2025 and 2026, make this urgent rather than aspirational. Regulatory pressure is the catalyst. Competitive advantage is the prize.

The firms that invest now - in data architecture, workflow automation, and governance frameworks - will define the next decade. Those waiting will find themselves unable to compete on speed, transparency, or cost.

Leadership teams ready to assess their operational foundations against these frameworks can contact Toucanberry to discuss readiness benchmarking and roadmap development.

> *The technology exists. The constraint is implementation.*

About Toucanberry

Toucanberry helps life reinsurers build operational infrastructure that makes AI valuable. We work with firms navigating AG55 compliance, BMA disclosure requirements, and the operational transformation needed to compete on speed and transparency.

Our approach starts with operations, not technology. We map workflows, measure cycle times, identify manual touchpoints, and build data architecture and governance frameworks that turn AI into a competitive advantage.

For more information: **www.toucanberry.com**

Research Methodology

This report synthesises regulatory guidance, industry research, and operational experience from Bermuda's life reinsurance market.

Operational and governance statistics draw from benchmark research by APQC, McKinsey & Company, Accenture, IBM, Deloitte, MIT and insurance trade publications, including Insurance Business and Reinsurance News. Regulatory requirements and timelines reflect publicly available guidance from the Bermuda Monetary Authority (BMA) and the National Association of Insurance Commissioners (NAIC).

Sources include:

- **Regulatory guidance:** BMA AI governance discussion paper, BMA enhanced disclosure consultation, NAIC AG55 adoption notice
- **Industry benchmarks:** McKinsey Global Finance Benchmark 2023, Accenture AI Governance Research 2024, IBM Global AI Adoption Index 2024, Deloitte AI Maturity Studies, MIT
- **Market research:** Conference insights from BILTIR 2025, Redirect 2025 and Singapore International Reinsurance Conference 2025
- **Operational case studies:** Operational case studies: Selected Bermuda life reinsurers (including Toucanberry client projects), Swiss Re ClaimsGenAI deployment, Aon Impact Forecasting, Zurich Insurance commercial automation

www.toucanberry.com