pwc

Reimagining the core of insurance

Driving intelligent decisions with smart systems

April 2026

Contents

01 Next-gen insurance 05

02 Connected core insurance systems 07

03 Data as the backbone of smart core systems 09

04 Data driven decision-making 12

05 Reliable GenAI at the core of insurance data 16

06 Real-world scenarios 20

07 A practical roadmap to modernisation 24

08 Next is now 27

09 Glossary 29

Foreword

Amit Roy
Partner and Leader, Insurance and Allied Businesses
PwC India

Insurance has always evolved through cycles of change, but rarely have the pressures converged with the intensity we see today. Insurers are being asked to move faster, decide smarter, and operate with greater accountability, all at once. Customers expect instant clarity, while regulators expect provable transparency and leadership teams navigate unprecedented complexity across legacy cores, fragmented data, and an expanding set of AI-driven solutions. Traditional approaches of rules locked in silos, opaque analytics, or ungoverned experimentation with GenAI will not suffice as the sector grapples with regulatory and technological changes.

This report explores how the insurance sector can evolve responsibly with the help of a decision-making framework with a smart core insurance system. By unifying trusted data, modular decision services, and GenAI within the core, the paper discusses how insurers can move beyond fragmented logic and episodic automation towards a cohesive, auditable decision fabric.

In each use case:

- A smart core system emits the event.
- Data mesh supplies trusted context.
- The decision framework orchestrates rules, models, GenAI, and human review.
- Outcomes feed back into the data mesh for continuous improvement.

Going beyond an architectural construct, what makes it distinctive is its insistence that speed and trust are not trade-offs. GenAI tools are implemented to summarise content, enhance decision-making and assist human agents within clearly defined guardrails. Confidence thresholds, lineage rich data, explicit fallback paths, and human oversight are embedded in the design to ensure that accountability is never an afterthought. In insurance, where every decision carries financial, regulatory, and reputational weight, this balance is essential. It reflects a mature view of GenAI adoption that aligns naturally with regulatory expectations while delivering tangible business outcomes.

The report also lays emphasis on a mindset shift: from isolated pilots to production platforms, from experimentation to institutional capability. The focus on telemetry, continuous feedback, and measurable outcomes underscores that responsible GenAI is not static but is learned, tuned, and governed over time. As insurers progress on this journey, the core itself evolves from a transaction engine to a trusted decision platform that strengthens customer relationships and reinforces confidence across the ecosystem.

The message is clear: the future of insurance belongs to organisations that pair innovation with governance, and treat decisions as a first-class, governed capability. This report can be considered to be a playbook for insurers who are ready to modernize their operations with trust at the core of the process. We hope this report inspires you to make informed decisions that are aligned to the market's pulse.

Foreword

Rajarshi Ghosh
Partner, Advisory
PwC India

The insurance industry, built on managing risk, is undergoing significant changes. In a time where customers want more while risks are changing, and regulations are getting stricter, what will really set us apart is how we use our data, not just for the past but as an active tool to shape our future. Valuable insights have been stuck in silos across organisations for many years which often impacts agility, accuracy, and the competitive edge of organisations. The way forward is clear—data is not just a byproduct of our operations, it is the key factor that drives every strategic decision, shapes every customer interaction, and supports every innovation.

This paper looks at how we can effectively harness this power, moving beyond isolated insights to a unified, smart future. Imagine having one reliable place where all the information is stored in an organised, easily accessible manner which enables businesses to recognise risks, offer personalised experiences to customers, and streamline their operations and shift from reacting to issues to making smart, informed plans. This central system, often built on platforms like modern smart core insurance system, ensures our data is trustworthy, follows rules, and is always available when needed.

However, with all this data comes great responsibility. The primary goal of this system is to use data to enhance customer service and trust by establishing clear rules for data sharing, control who can access it, and ensure every piece of data is accurate, traceable, and used fairly. GenAI can be instrumental in this with its ability to quickly read long documents, summarise notes from adjusters, write clear explanations for customers, and instantly inform the organisation about any changes in an account. A smart system could then bring these insights directly into the tools for underwriters and adjusters and provide clear reasons and confidence levels for its suggestions.

However, these systems still keep human review as the central point of control along with documenting every decision. This paper discusses how we can create this data-driven ecosystem. We invite you to explore these key ideas with us as we work together to build a smarter, more trustworthy insurance future, truly powered by data.

01

Next-gen insurance

Illustrative

I filed a claim after a minor accident, but I keep repeating my information every time I call.

Polly, a policyholder, expects an effortless claims experience after a minor accident. However, instead of a seamless experience, she faces long wait times, repetitive form filling, and unclear communication.

(Representational Image)

I juggle several outdated systems that don't talk to each other. Manual tasks slow me down and make handling exceptions difficult.

Across town, Mike, a claims adjuster, juggles multiple systems with outdated interfaces, struggling to access timely information to serve customers effectively.

(Representational Image)

Millions of people depend on insurance for safeguarding their homes, vehicles, and businesses. However, insurance claims processes are often not so smooth.

Consider the example of Polly, a policyholder, who wanted to update her policy online and expected the update to take only a few minutes. Instead, the process stretched longer than expected as she had to enter the same details more than once and wait for updates from different teams which made the process confusing and slow.

Mike works as a claims adjuster. His day involves managing several cases at once where he reviews documents, checks policy details, and coordinates with different teams. For him, the most time-consuming process is gathering information.

Sarah is an underwriter. Her job depends on experience and judgement. However, for her, the hardest part is not the analysis but collecting the information needed to make that decision which sometimes takes longer than the decision itself.

For all these people, legacy systems slow them down when customers increasingly demand faster resolution of their queries. The employees too want to spend less time on administrative tasks and focus on other problems.

At the centre of these challenges lies one thing: **Decisions**. Insurance operations depend on thousands of decisions every day, including:

- Who should be insured
- At what price
- Which claims should be approved
- And when should a case be escalated?

Some domains of insurance move very quickly. Auto and home insurance, for example, deal with high volumes, making speed and consistency the primary factor in these domains while commercial or specialty policies require deeper analysis and expert judgment instead of faster resolution.

Faster decisions do not mean careless decisions. The first step of smart decision making is simple: insurers need one reliable view of information. Without that, teams end up repeating the same work repeatedly.

Good decisions begin with reliable information. Everyone should work from the same source of data. This includes customers, policies, claims, and risk details. Rules need to be clear, and they should be easy for teams and partners to access. Not every case needs the same level of attention. Simple and low-risk cases can move automatically. More complex situations should be managed by experienced professionals who can review them with the right context.

Decisions also need clear guardrails. Every important decision should leave a record. This helps maintain fairness and also supports regulatory compliance. Improvement of processes should be continuous, and feedback from past decisions could help teams refine their data, update the rules, and improve training over time. The systems supporting these decisions must be dependable. They need to remain stable during daily operations while remaining flexible. Teams should be able to communicate seamlessly, and information should move quickly across different channels. This is where a smart core system, which provides stability and flexibility to adapt to various scenarios, becomes essential for organisations.

A data mesh can help teams access trusted data more easily and a decision framework helps apply consistent decision logic across the organisation which enables them to make faster and more reliable decisions.

Insurance is not just about policies or transactions. It matters most in uncertain moments where people need answers and reassurance.

The insurance industry has a clear opportunity to transform itself by reimagining how people, information, and processes meet to create faster, fairer, and more transparent decisions. This approach can lead to greater operational efficiency and better customer outcomes in an increasingly complex and challenging environment.

The challenge is not simply improving systems but enabling faster and more intelligent decisions across the entire insurance value chain. Chapter 2 introduces how a smart core system, combined with a decision framework approach, can transform fragmented processes into a connected decision ecosystem.

Insurance doesn't fail from a lack of effort—but from fragmented and slow processes. A smart core system, combined with a decision framework could unlock a unified, intelligent decision system that helps organisations make faster decisions.

02

Connected core insurance systems

Unifying smart core system, decision framework, and data mesh to shape modern insurance

Sarah, an underwriter, has to deal with a lot of siloed data and usually relies on her experience and judgment to assess insurance risks. However, it is cumbersome to consolidate the necessary information from disparate systems before she can make informed decisions. This not only increases the turnaround time but could also lead to errors, inconsistencies, and delays in policy issuance.

A **smart core system** could help Sarah by providing a unified, stable core system that integrates data from multiple sources into a single view. It can also automate routine tasks, allowing her to focus on analysis and decision-making.

How decision framework and data mesh complement smart core systems

In today's fast evolving world, insurance landscape is undergoing tectonic shifts due to rising customer demands and changing regulations. These necessitate changing and expanding the foundational infrastructure by integrating decision framework and data mesh architectures that complement the smart core system.

A decision framework is a structured approach that analyses data to drive decisions. It follows a decentralised approach, ensuring agility and scalability while maintaining governance through centralised oversight embedded within the smart core system.

Modernising insurance with smart core systems

Modernisation doesn't mean replacing older, outdated legacy systems, but building future-proof core systems that can seamlessly be integrated with existing infrastructure, streamline workflows, and improve operational flexibility while enabling innovation.

Smart core systems represent new generation, unified platforms designed with the capability to automate decision making while maintaining governance, regulations, and managing emerging market demands. By integrating smart, automated decision-making features within a regulated and compliant framework, smart core systems enable insurers to expedite core operations while maintaining accuracy and control.

Data mesh is a decentralised approach to data ownership designed to address common challenges like siloed data, delayed real-time access, and governance complexities. It is based on **4 core principles**:

01

Domain oriented data management: Data ownership and responsibilities resides with specific business domains who are well-versed with and closest to the data.

02

Data as a product: Data is not treated as a by-product but as a product with well-defined ownership, quality and usability.

03

Self-serving data infrastructure as a platform: User-friendly data platform that can be accessed and managed independently.

02

Federated data governance: A hybrid approach that follows centralised standards with decentralised execution.

Collectively, these principles ensure higher standards of quality, interoperability and compliance with respect to data

At the core of this entire framework lies a critical enabler: data. Relevant, consistent, and quality data serves as the primary raw material for driving efficient processes.

03

Data as the backbone of smart core systems

Governed data as the driving force for powering smart core systems, data mesh, and decision framework

Insurance industry relies heavily on data, but ungoverned siloed data makes the underwriting and claim process challenging. Claim settlement is one such area which heavily depends on different data both internal and external. A centralised data lake system feeds data to a smart core insurance system, which helps to develop a single source of truth. In a world where every decision is data driven, accurate and secure data optimises the **AI model's** performance and helps them to make fast, consistent, and informed decisions.

Core systems comprising policy, billing, pricing, underwriting, and insurance centres, now feed data to smart insurance core system. These data products interchange data on a real-time basis to generate accurate results. However, data interchange without proper data governance can generate inaccurate results.

Data contracts provide data governance and act like ongoing agreements to spot mistakes, changes in data structure, and any compliance problems using continuous and regular GenAI validation and anomaly detection.

Smart core systems, built on modern insurance core platforms, follow an architecture where several processes work together in a coordinated manner. Data moves from source systems (policy administration, claim intake and adjudication, payment processing, etc.) into a central repository system through automated pipelines. For seamless coordination between enterprise workflow orchestration tools, and smart core insurance platform services integrate data across domains. GenAI models detect failures early and fix them to ensure that policy updates, claims decisions, and billing transactions stay in sync. **Metadata governance** uses open, industry-standard schema definitions and data contracts, while GenAI models keep an eye out for any unexpected changes in the data structure. Any unauthorised or accidental changes can be quickly detected by smart core systems to keep data structures safe and consistent across different areas.

AI-powered observability detects anomalies, pinpoints root causes, and alerts engineers before issues reach business users, reducing downstream impact, speeding up issue resolution, and enabling comprehensive end-to-end data quality monitoring.

Access to data can be controlled based on people's roles, using systems like active directory and access management system to manage who can see what. Special tools monitor how people use their access to catch anything unusual or suspicious early on. Protecting private information (like personal details) is very important and can be done by restricting access or anonymising data in different ways. The system follows important privacy laws like GDPR, HIPAA, and PCI DSS, and keeps clear, unchangeable records to ensure data security. Every time data is accessed or changed, a record is made and sent to monitoring tools. These tools use smart technology to spot problems before they happen, so the system stays safe and runs smoothly.

From underwriters' risk modelling in the policy management system to fraud detection in the claims management system, and revenue assurance in the billing management system, the AI-enabled data ecosystem operates as a resilient, self-optimising fabric. Teams govern confidently, adapt quickly, and innovate continuously, thereby improving processes across the insurance value chain.

Illustrative

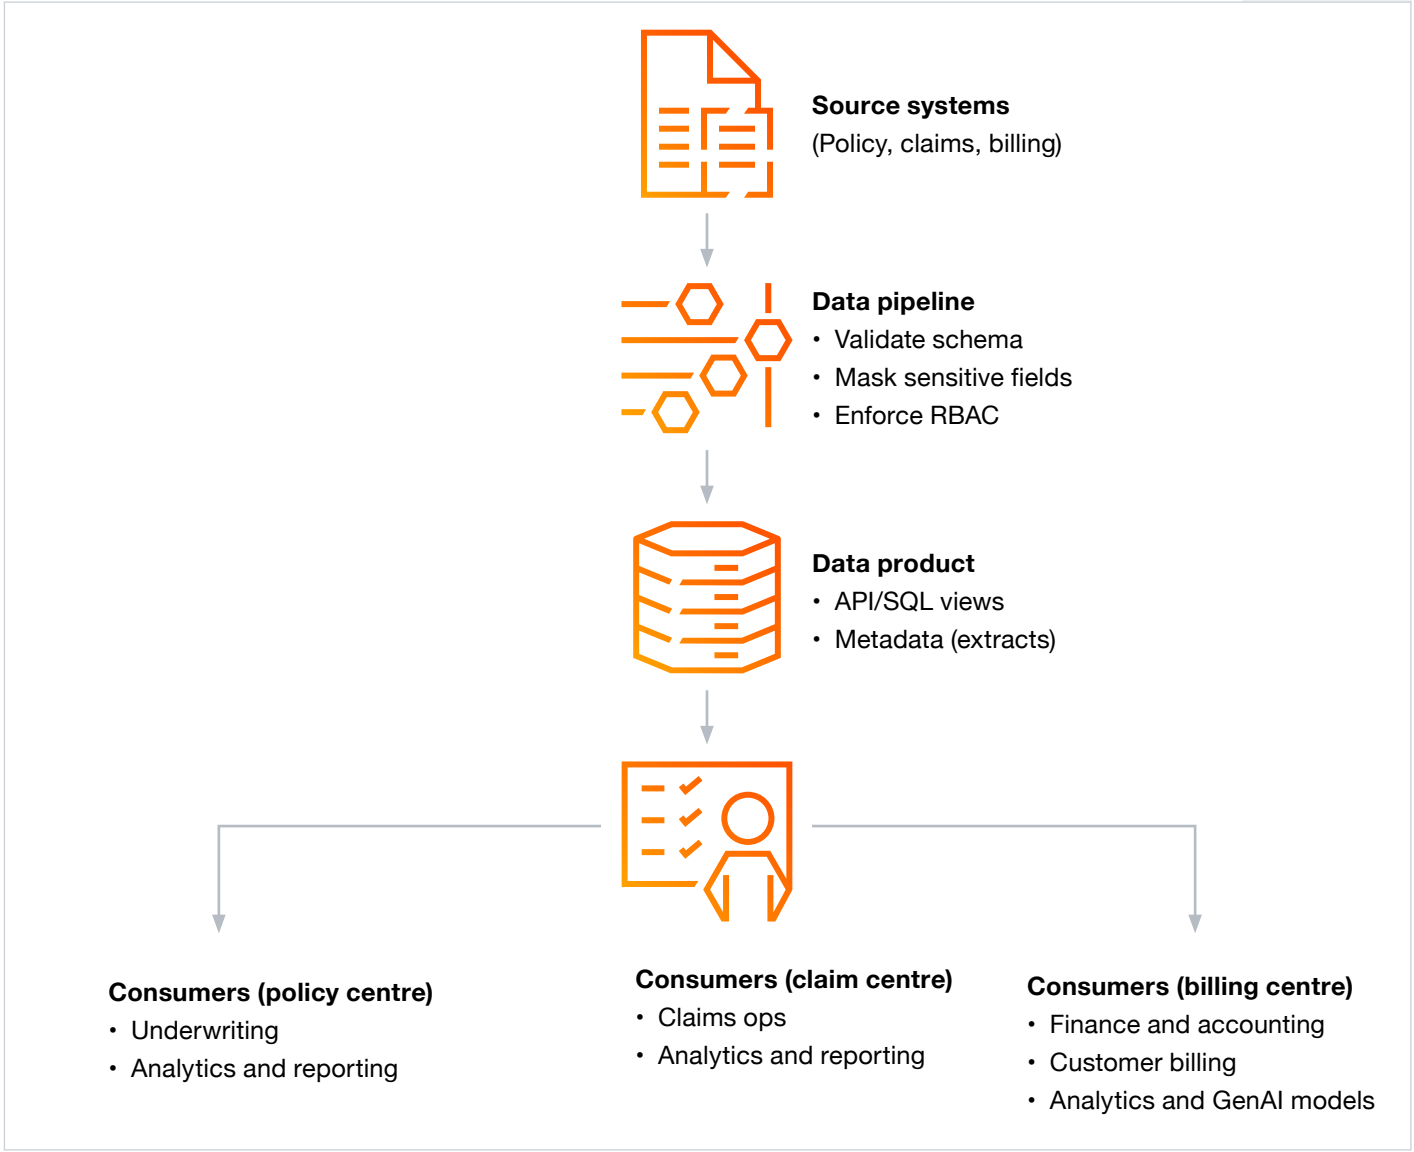

Attribute	What it includes	Key details	GenAI impact
Schema	Strict typing; JSON; versioned	Controlled change; full lineage	Consistent inputs; fewer errors/bias
Freshness	Near real-time policy; hourly claims; daily billing	SLOs monitored and alerted	Timely insights at decision time
Ownership	Domain teams (Policy, Claims, Billing)	Quality stewardship; issue management	Accountable, ethical data
Access	RBAC + SSO (AD)	Least-privilege; regional controls	Safe, compliant GenAI operations
Privacy	Masking/tokenisation/anonymisation/synthetic	At rest and in transit	Protects PII; enables sharing/testing
Observability	Central logs/metrics/traces; anomaly RCA	Predictive risk alerts; self-healing	Reliability before impact
Compliance	GDPR/HIPAA/PCI DSS	Immutable audit logs; automated checks	Regulatory trust and transparency

A data product is an end-to-end managed offering defined by its elements—ownership, schema, quality, access, privacy, observability, and compliance—but nondata stakeholders often struggle to understand this metadata for a given dataset. Data product elements define the full managed offering which can be difficult for a non-data background stakeholder to understand with the metadata of a specific dataset. To bridge this gap, a data contract card is considered. A data contact card is a quick, one-page summary that shows important details about a specific set of data. The card outlines access policies controlling data visibility, compliance badges reflecting regulatory adherence, and the current validation status which shows whether the data has been validated or flagged for anomalies, along with the latest validation timestamp.

This centralised view enables data consumers and stewards to quickly evaluate data trustworthiness, accessibility, and compliance, and support informed decisions and strong data governance.

Data contact card (claims)

Mike
Claim adjuster
Schema v2.4.1
(Representational Image)

✓ **Validated**

Freshness SLO
SLA < 15 million | Actual: 12 million

Key fields
- ClaimID
- PolicyID
- Loss date
- Cause of loss code
- Location
- Coverage status
- Reserve amount
- Subrogation potential score
- SIUFlag

Role	Access level
CEO	Full access
COO	Full access
CDO	Governed access
Claims manager	Write
Claim adjuster	Write
Underwriter	Limited write
Analyst	Read only
Customer service	Restricted
External partner	No access

Privacy and protection

PII and PHI sensitive

HIPAA-regulated

Masked and audited

Compliance badges: HIPAA, GDPR, SOC
Last checked: 3 April 2026 | 10:30am UTC

Data contact card for claims, showing ownership, key data fields, access control and compliance status in single view

04

Data-driven decision-making

Building trust through responsible actions

As an event happens, the first step is the immediate identification of the event and recording of all the facts in detail, leaving no stone unturned. This is immediately communicated to the stakeholders involved in the process. Afterward, a process of filtration and prioritisation of the events starts, concentrating only on what is of utmost importance. Once a thorough analysis of the event is done, its context and implications are clearly understood, a course of action is determined and implemented with the utmost speed and vigilance.

The entire process of policy lifecycle—including customer onboarding, underwriting, and binding—results in events being emitted by the smart core system, which are further enriched by the data mesh. GenAI validates the input, simulates risks, and generates clear communications, whereas the decision framework provides recommendations regarding the next line of action, including complete and approve, price and bind, refer and human review. Similarly, in the claim lifecycle process, reporting and intake, assessment and investigation, and settlement and resolution, FNOL, and evidence updates are parsed by GenAI. It also extracts and categorises the data, detect fraud, predict outcomes, and determine the line of action, including scheduling an inspection or settlement range.

Application of GenAI across policy and claim lifecycle

Customer onboarding

The information starts to flow into the data mesh as soon as Polly submits her case in smart core system, which validates details against internal records and trusted external sources. The decision framework asks Polly clear follow-up questions based on gaps or anomalies in the given information and also flags potential risk and, if needed, re-routes the case to Mike for review. Next best step for Mike is to either complete-and-approve or escalate with appropriate reason.

Claim reporting and intake

Polly files FNOL in smart core system; the data mesh extracts and validates data from photos, videos, and forms. Severity check, fraud scan, and selection of the right queue are done by the data mesh itself, and then it schedules the first action—such as an inspection or estimation. Next best step is starting an investigation with the prioritised tasks.

Policy underwriting

Via smart core system, an underwriting event contacts Sarah with governed data from the data mesh. The decision framework simulates claim scenarios, and with confidence and fairness checks it recommends a price range and explains key risk drivers. In case thresholds are exceeded, it prompts human intervention. Next best step is price-and-bind or refer for high-level assessment.

Claim assessment and investigation

When information comes via the smart core system, the data mesh organises policy conditions, previous claims, and repair or medical documentation. The decision framework points out inconsistencies, proposes an investigation strategy, and forecasts potential results. It advises whether to approve, deny, or escalate decisions, along with justifications.

Policy binding

Once approved in the smart core system, the data mesh is triggered by the binding event to publish the finalised terms. The decision framework writes a personalised, simple summary for Polly and guides her through all the documents, inspections, and payment. For audit purposes, confirmations are captured. Next best step is activating coverage and hand off to servicing.

Claim resolution and settlement

When assessment concludes, the smart core system emits the decision event and the data mesh records outcomes. The decision framework proposes a settlement range and triggers a conversational assistant to explain options, collect any missing items, and confirm acceptance. Feedback loops update models and rules. Next best step is issue payment or negotiate alternatives.

Imagine Polly files a claim due to hailstorm damage. Let's walk through how the process advances and how GenAI helps make it efficient.

The first step is filing the claim, where the system collects basic information—Polly's policy details, coverage, and date of loss. The system quickly verifies if Polly's policy was active at the time and if hailstorm damage is covered at her location. This early check helps decide if the claim should move forward. GenAI speeds up this step by reviewing policy documents, highlighting relevant coverage, and summarising key points for claims handlers.

Next, the claim is checked for any policy exclusions, and the likely severity of damage is estimated. GenAI simplifies complex policy and claim info into clear summaries for the claims team. It also helps keep customers informed about the process and factors like deductibles.

Risk assessment in the workflow ensures high-risk claims receive proper approval, while low-risk claims are processed faster—balancing speed and control. All claim data is securely stored for future analysis, improving accuracy and transparency.

Technologies like GenAI, combined with a solid decision framework, transform insurance claims and policy management. These innovations reduce manual work and increase decision quality by enabling faster, clearer information handling. By automating routine tasks while keeping human oversight where needed, insurers deliver a smoother, more transparent experience for customers like Polly, while strengthening compliance and operational strength.

Insurers are increasingly being encouraged to invest in these smart solutions as part of their digital transformation to stay competitive and drive lasting business value.

As we advance into these decision-making technologies, it is equally important to ensure security during every step—this facilitates fast and efficient processes that are also resilient and trustworthy.

New claim, suspected hail

Illustrative

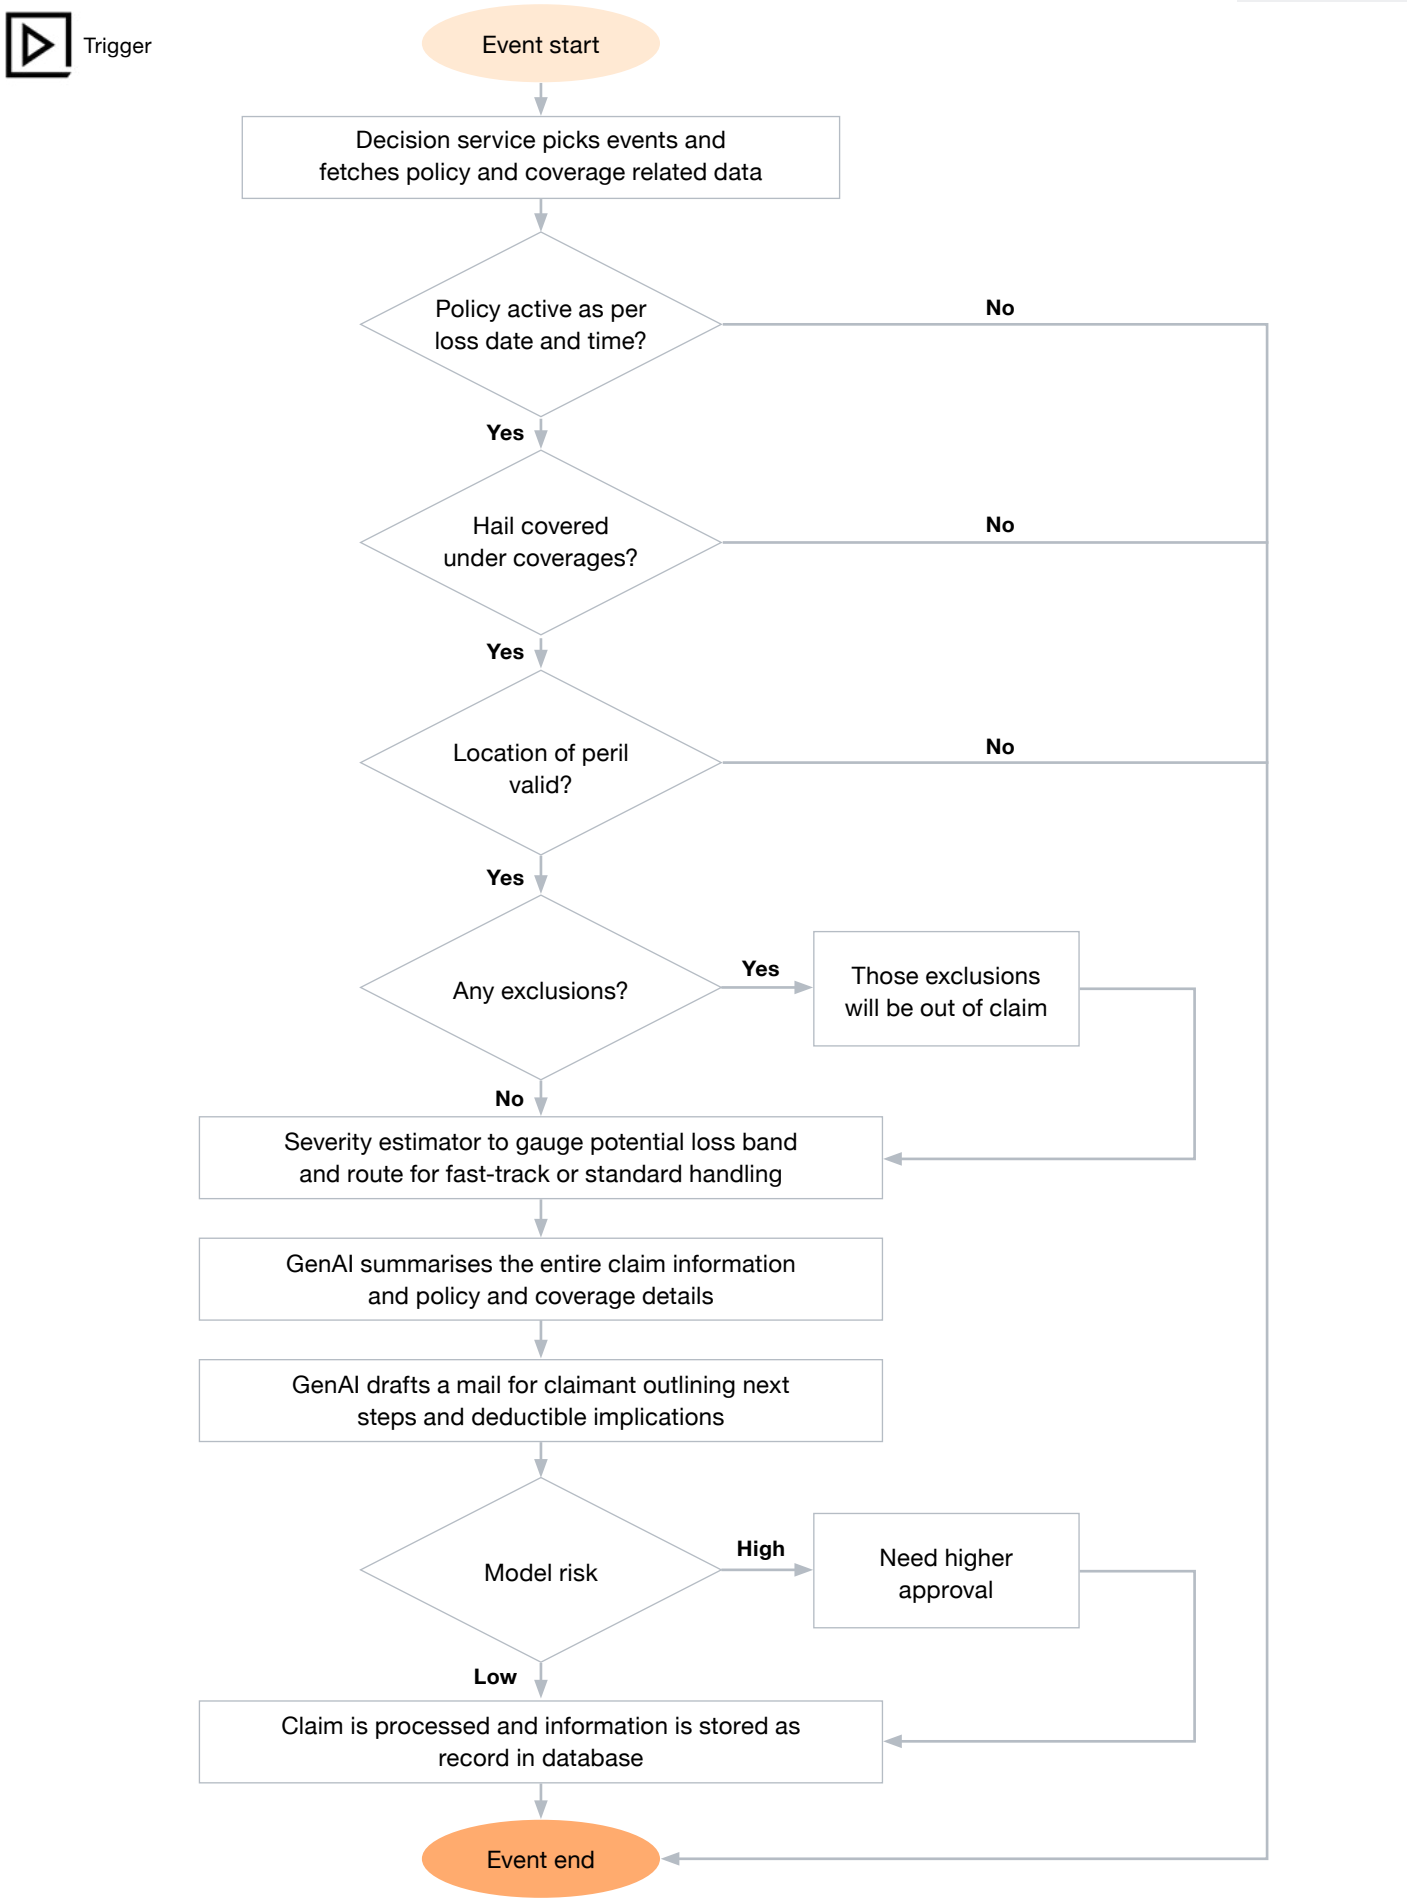

05

Reliable GenAI at the core of insurance data

Secure data, transparent logic, controlled actions

Let's consider a scenario where you are at a busy airport. There is a smart control tower that guides all the planes by telling them where to land, when to take off, and which runway to use. Now, picture that the control tower started giving directions based on guesses or unclear information. This may result in serious accidents as unclear instructions will not facilitate a smooth landing. This is like how GenAI works in insurance. When agents or customers ask questions about policies, GenAI provides quick answers. But if the GenAI system isn't carefully controlled, it might give confident responses based on incorrect or outdated information. This can confuse customers and create compliance risks for the company.

To avoid these problems, insurance GenAI must work like a well-managed airport control system, relying on only trusted data such as official policy records and verified endorsements. Every response should be traceable, with a clear record of where the information came from, so any answer can be checked and verified.

Illustrative

The question
Friday afternoon; agent asks GenAI, "Is flood damage covered for this policy?"

The reply
Instant, confident—but wrong

The cause
Generic template used; endorsement ignored; clause superseded months ago; no citation

The fall-out
Agent relays answer; claim filed; two weeks later underwriting flags it—coverage never applied

The lesson
Compliance breach, customer complaint, manual audit. No malice or outage—just GenAI without guardrails.

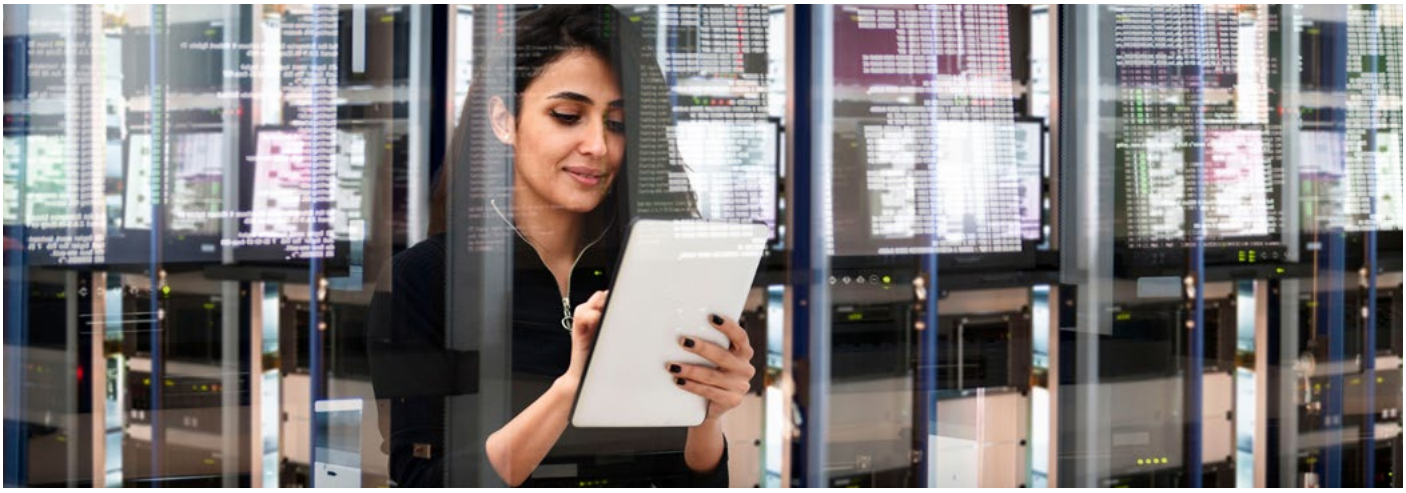

Protecting personal data is just as important. Just like airport security limits people who can enter secure areas, insurance companies use strict **access controls** to keep sensitive customer information safe.

In practice, this means putting strong guardrails around the GenAI's answers:

GenAI uses carefully worded checklists, like **prompt templates** for questions, to ensure it follows the safety rules and cites where it got the information.

Dynamic filters, like security scanners, block any unsafe or wrong answers before they reach customers.

If the system is unsure about an answer, it follows a **fallback path** and simply says, 'I don't know', instead of making a dangerous guess.

By combining a smart core system with a solid decision framework and reliable data sources, insurance GenAI becomes fast, safe, and trustworthy. Every answer is linked to real company data—easy to explain and fully auditable.

Prompt templates

Prompt templates are prewritten instructions that tell the GenAI system exactly what to use, how to reason, and how to present the answer. In an enterprise smart core system setup, they go through the decision framework, which pulls only governed facts from the data mesh (including policy records, endorsement text, and effective dates). The reply is then formatted with a clear decision, a short rationale, citations, and a confidence score while refusing to guess if any required fields are missing. For example, when an agent asks, "Is flood damage covered?", the template requires the policy ID, retrieves the active endorsement and dates from the enterprise smart core system via the decision framework, and returns a simple 'yes' or 'no', along with the exact clause and timestamp. This matters because it turns every response into a repeatable, auditable step that reduces errors, speeds up reviews, and provides customers and regulators with a clear view of the source.

If an agent asks about flood coverage on a certain policy number, the prompt template ensures the GenAI checks only that policy's details. It gives a clear answer based on the latest endorsement, along with the exact contract clause it finds.

Dynamic filters

Dynamic filters are automated guardrails that scan inputs and outputs for risk before anything reaches the user. They are present in the decision framework and use modern insurance platform rules and data mesh metadata to check for missing citations, PII leakage, toxic language, and conflicts with coverage rules, locations, perils, and waiting periods. These filters can block, redact, or downgrade a reply and log the reason. For example, if the model states that hail is covered but the location-peril combination is excluded or a named-storm waiting period applies, the filter stops the message, triggers a safe re-retrieve of the correct endorsement, or asks a clarifying question. This matters because it catches problems early, protects privacy, and keeps communications compliant without slowing down daily work.

If the GenAI system says, “Hail damage is covered”, but the policy is for Texas where hail is excluded, the filter will intercept this answer, update it with the correct information, or ask for more details before responding.

Fallback paths

Fallback paths are controlled routes the system uses when confidence is low or sources disagree. This mechanism prioritises a careful pause over a wrong answer. The decision framework monitors grounding signals and thresholds, then decides whether to re-query enterprise smart core system records through the data mesh, run a rules engine check, ask the user for a missing detail, or hand off to a human reviewer—while logging each step for audit. For example, if the endorsement effective date is unclear, the GenAI system responds with: “I don’t know yet; please confirm the policy number or date.” It then pulls the right record and returns a cited decision or escalates to an expert. This matters because it prevents compliance breaches, avoids rework and complaints, and builds long-term trust by making every decision explainable and recoverable.

If an agent asks about coverage but the GenAI system is not sure which policy version is active, it will ask, “Please confirm your policy number or effective date,” then check again or escalate to a human if needed.

With trusted guidance and diligent controls, reliable GenAI is built on accurate data, transparent decisions, and strong safeguards—enabling people and GenAI to work together more effectively across a wide range of scenarios.

Guardrails turn answers into evidence

Before

Coverage:
Flood is covered
- No citation
- No effective date
- No confidence

Prompt templates
Structured cited answers

Dynamic filters
Block unsafe outputs

Fallback path
Safe recoveries on low confidence

After

Coverage: Not covered
- Citation: Endorsement
- Effective date: 25 January 2024
- Confidence: 0.92 (threshold 0.85)

  – Citations present: 0% → 100%
  – Confidence available: No → Yes (0.92)
  – Compliance risk: High → Low

Safety checklist

Use only trusted data products for input.

Scan responses for privacy issues, bias, or harmful content.

Limit access to sensitive information to authorised users only.

Auto-approve answers only when confidence is 85% or higher.

Ensure AI answers come from verified, trusted sources.

Require human review for important or risky decisions.

06

Real-world scenarios

Powered by clean data, smart systems, and human judgment

When insurers set responsible GenAI on top of modern smart core insurance platform and governed data mesh, the benefits aren't theoretical—they show up in the numbers. Faster cycle times, fewer errors, and decisions that stand up to audit. The magic isn't in one big leap; it's in a series of high-impact use cases where GenAI does the heavy lifting and humans stay in control.

Each use case follows a simple pattern:

- **Inputs:** Clean, trusted data from modern insurance core platform products and curated partners
- **Actions:** Rules fire, models score, GenAI drafts with citations
- **Outputs:** Clear summaries, checklists, and next steps—ready for human approval
- **Business impact:** Measurable gains in speed, accuracy, and governance

From compressing FNOL intake from hours to minutes, to reducing underwriting errors and accelerating subrogation prep, these wins tie directly to KPIs insurers care about: cycle time, cost per decision, and compliance confidence.

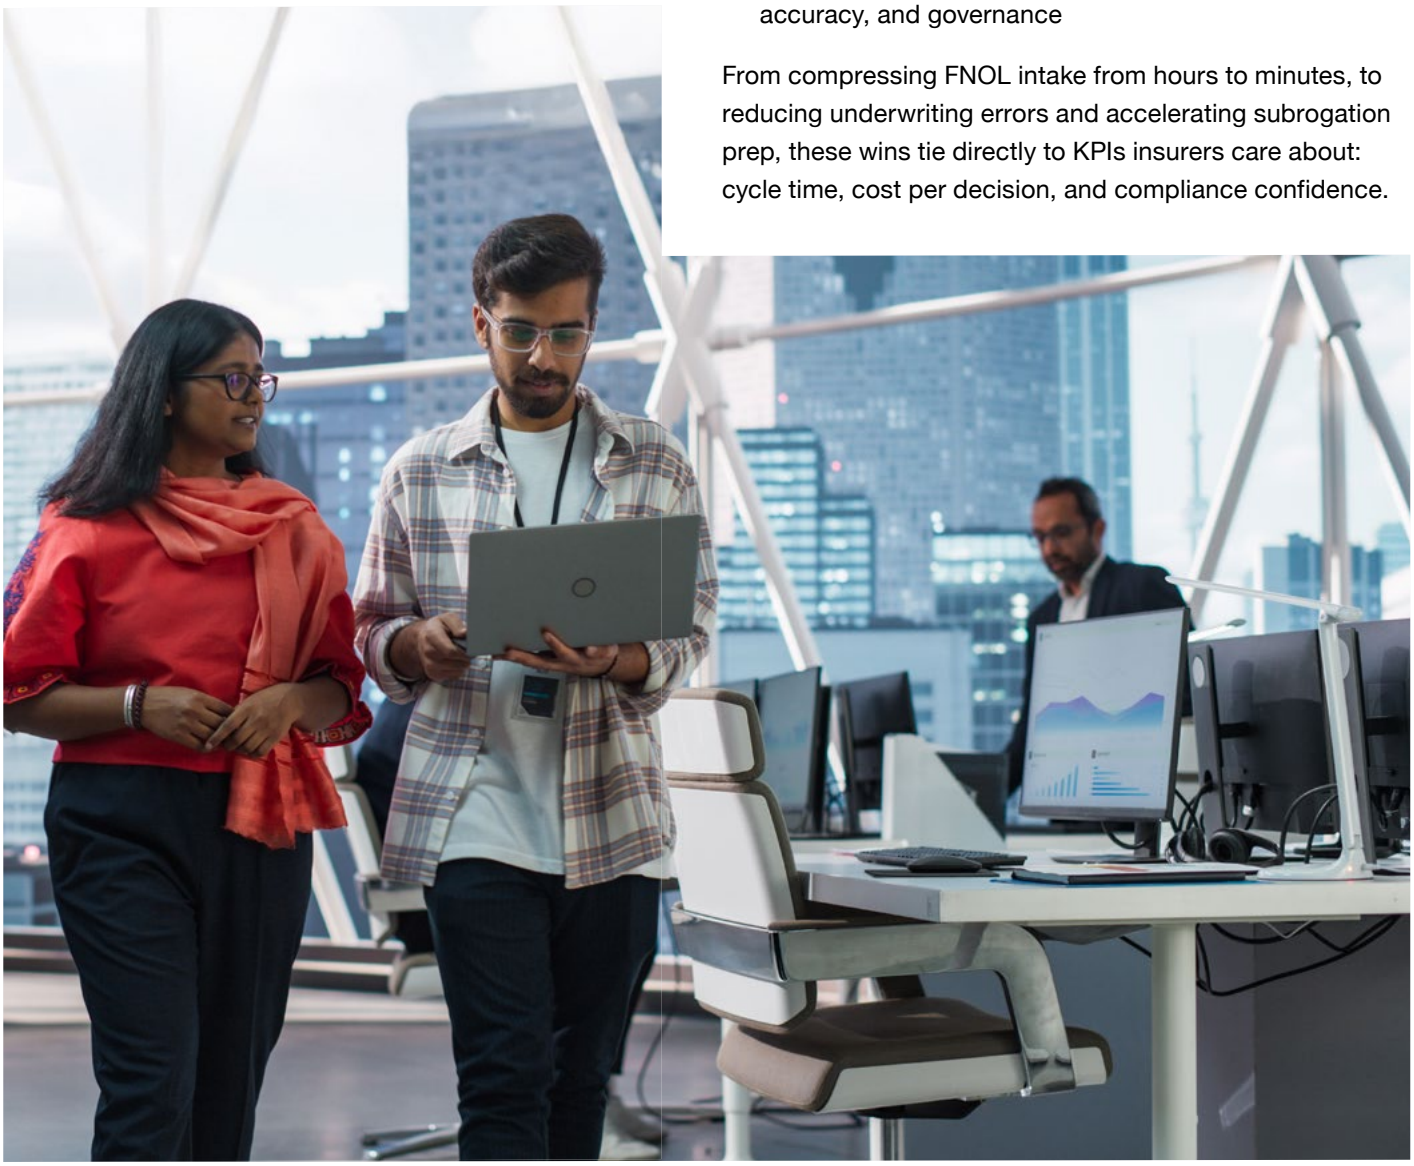

Scenario 1

FNOL summary

8:45am

Polly submits a first notice of loss after a minor vehicle collision.

The information arrives in pieces mobile photos, short voice notes, and existing policy data already stored in the core claims management system.

Instead of forcing the adjuster to assemble the story manually, the decision framework activates governed by trusted smart core insurance data products, orchestrated by responsible GenAI.

Evidence ingested and interpreted

Six photos analysed:

- Bumper and headlamp damage detected, no structural frame impact identified

Voice note:

- Automatically transcribed, key facts mapped to loss description

Structured data:

- Policy and coverage pulled from the claims centre

8:50am

Prior claims check completed (last three years: none). Each insight is linked to its source—image metadata, voice transcript, or policy clause-ensuring full explainability.

Guided decision checklist (decision framework in action)

- Confirm repair network eligibility.
- Validate vehicle registration vs policy.
- Auto-segment claim to fast-track lane.
- Trigger digital repair estimate request.

The checklist is not static, but it is dynamically generated based on policy rules, claim context, and confidence thresholds—all governed within the decision framework.

Claimant communication (drafted, not decided) pre-approved message shown to the adjuster.

8:57am

“We’ve reviewed your policy and initial loss details. Based on the information provided, your claim qualifies for fast-track processing.

Your deductible is ₹5,000.

Please upload any additional photos if requested to proceed without delay.”

Underwriting intake summary

Scenario 2

10:15am

Sarah, a commercial underwriter, opens a renewal submission from a broker. The package is dense ACORD forms, spreadsheets, loss runs, engineering reports, and over **200 pages** of attachments. In the past, this would mean days of manual review, cross-checking appetite, and chasing missing details.

This time, the decision framework activates across governed policy management system data products, powered by responsible GenAI.

10:40am

Documents parsed and insights extracted

- ACORD forms

Locations: 5 manufacturing sites

Construction type: Mixed (two masonry, three-metal frame)

- Spreadsheets

Annual revenue: 7420 crore
Payroll exposure: 765 crore

- Attachments (200+ pages)

Engineering report highlights summarised
Fire suppression compliance confirmed at 4 out of 5 locations
Gap identified: Updated roof inspection missing for location #2
Guided underwriting checklist (decision framework in action)

- Risk validation completed
- Loss history compared against peer benchmarks
- Engineering recommendations reviewed

11am

Broker communication (drafted, not dispatched)
Pre-drafted message for underwriter approval:

"We've completed an initial review of the renewal submission.

The risk remains within underwriting appetite.

To finalise terms, please provide the latest roof inspection report for Location #2.

Once received, we can proceed without delay."

Human approval remains mandatory-responsible GenAI supports decisions, it does not replace them.

Subrogation evidence pack

Scenario 3

10:15am

Renshaw, a subrogation specialist, opens a recently settled to claim. He tracks down police reports, reviewing liability notes and searching for third-party coverage details scattered across systems and attachments.

This time, the **decision framework** activates automatically once the claim reserve is confirmed.

Evidence assembled and prioritised

Police report

Fault assigned to a third party driver

Timestamp and location validated

Liability notes

Adjuster observations summarised

Impact direction supports adverse liability

Third-party coverage

Policy number and limits extracted

Coverage active on loss date

Carrier contact details verified

All documents and insights are traceable to their original sources, with page-level citations preserved.

Subrogation readiness checklist (decision framework in action)

- Reserve threshold met for subrogation
- Liability confidence above 70%
- Evidence completeness score: 92%

Post-settlement intelligence (beyond assembly)

After settlement, responsible GenAI scans:

Claims centre reserves and payment history

Third-party documents via CCC Safekeep integration

Prior recovery outcomes for similar claims

Using ML-based scoring, the decision framework:

- Ranks evidence by legal strength
- Highlights gaps that could weaken recovery

Renshaw doesn't start from scratch but he starts with a ranked, actionable dossier.

Scenario 4

Role aware service bot

Polly, a policyholder, opens the insurer's digital portal to check on her auto claim. At the same time, inside the contact centre, a service agent handles back to back coverage and billing questions. Traditionally, both journeys would involve queues, transfers, and inconsistent answers.

This time, the role-aware service bot activates—powered by responsible GenAI and governed by smart core insurance data products. Access is role-scoped. No overexposure. No guesswork. Two users, one question, and two right answers.

Polly : What's the status of my claim?

Bot response: Your claim is currently under review. The estimated completion date is 28 July 2026. No additional documents are required currently.

Agent response (operational detail):

- Claim status: Evaluation in progress
- Assigned adjuster: Mike
- Pending activity: Repair estimate review
- SLA: On track (T+3 days remaining)

Guided escalation (decision framework in action):

- Tier1 query resolved by bot
- Data confidence score above threshold
- Billing dispute detected
- Escalated to human agent
- Context packet attached automatically

When complexity rises, the decision framework routes the interaction—with the full context preserved, not a cold handoff.

Why it matters?

Use case	Before	After
FNOL summary	Manual and fragmented information handling with high effort	Automated summarisation with improved accuracy and faster claim readiness
Underwriting intake	Lengthy manual reviews with inconsistent assessments	Streamlined intake with faster decisions and improved consistency
Subrogation evidence pack	Time consuming evidence collection across multiple sources	Faster evidence assembly with better recovery preparedness
Service bot	High dependency on contact centres and frequent escalations	Improved self service, reduced agent workload, and better customer experience

Effective planning and strategic implementation of use cases serve as the cornerstone for driving genuine transformation within the insurance industry.

07

A practical roadmap to modernisation

Plan strategically, act decisively, and lead transformation

To successfully launch the smart core system implementation, we need to first clearly define the pilot project by focusing on one specific area such as homeowners FNOL coverage summaries.

To illustrate the potential impact of this approach, the examples and comparisons in this section are indicative and focus on relative improvements.

Secondly, we need to set a pilot period of about 8 to 12 weeks, with well-defined goals like reducing manual overrides, maintaining strong safety standards.

Next step involves building a small team with the right expertise and skillset, including claims and underwriting specialists, data stewards, platform and machine learning engineers, a GenAI risk and prompt lead, and a product owner to guide this whole process.

Thereafter, we'll use the smart core system to provide the required event data, while leveraging data mesh for clean, governed data. Apply rules, models, and GenAI through the decision framework system, making sure all actions are tracked for transparency.

Before going live, gather baseline measurements to understand current performance, and run tests to validate the new process. Begin the rollout in shadow mode—where the system processes data but doesn't affect decisions. Finally, move to a small live test once the desired results are reached.

Operations should be run on a weekly basis to review key metrics and cases, prioritising fixes in data, prompts, and models, supported by AI-driven incident management for quick detection and resolution. Clear roles ensure smooth delivery, compliance, and feedback, while expansion happens gradually with strong controls and transparent monitoring.

What to measure (AI-analysed continuously)

- Track cycle time, accuracy, override rate, cost per decision, and user satisfaction.
- Monitor safety signals, escalations, blocked prompts, and policy deviations.
- Use GenAI to spot patterns and drift, classify errors, and suggest the next tuning actions.

How to roll out safely

- Begin in shadow mode with drafts only and compare with business-as-usual for 1–2 weeks.
- Move to a canary release (10–20% of events) once metrics are stable.
- Keep automated rollback and kill switches for the GenAI step and the full flow.
- Rely on AI-assisted governance to watch KPI thresholds and trigger guardrails.

Keep cost under control

- Cache common lookups (e.g. weather, peril mappings).
- Keep prompts short and pass structured fields with citations instead of long text.
- Use lightweight classifiers for routing, heavier LLMs only for complex drafting.
- Fetch only the data you need from each product.

Pilot dashboard

22s
Time to decision

94%
Accuracy

12%
Override rate

$1.50
Cost per decision

Safety incidents

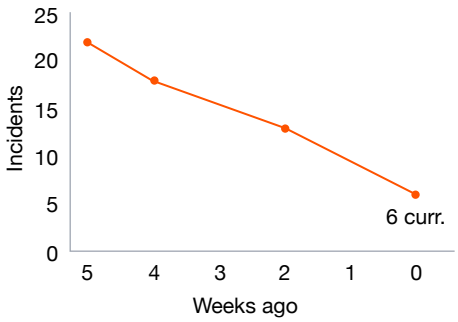

User satisfaction

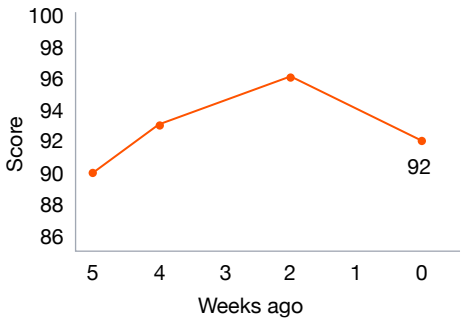

GenAI monitoring

Dimension	Indicator
Fairness	Stable ✓
Model drift	Moderate (6)
Prompt safety	8 blocked
Policy risk	2 deviations

Top override reasons and fixes

Rank	Reason	Fix in progress
1	Ambiguous input data	Improving validation logic
2	Lack of training data	Expanding datasets
3	Sensor malfunction	Hardware diagnostics
4	Complex scenario	Enhancing algorithms

Fairness/drift

Signal type	Indicator
Escalations	▲ Increasing
Blocked prompts	Active
Policy deviations	Detected
Recommendation	⚠ Satisfaction ⬇

The project progresses through four key phases, beginning with assessing and preparing governance, success metrics, and system integrations. This is followed by piloting and building core components alongside GenAI workflows. Next, pilot results are validated, models optimised, and governance protocols— including transparency, bias mitigation, and human oversight— are ensured. Finally, the solution is scaled enterprise-wide with ongoing training, adoption strategies, and continuous improvement.

This holistic roadmap starts with a focused, value-driven pilot targeting a specific business journey executed over 8–12 weeks, emphasising **decision speed**, **override reduction**, and **safety**. A cross-functional team leverages smart core system events, governed data mesh assets, and decision framework intelligence to ensure accuracy and responsible automation.20 Following successful shadow-mode testing, a small canary release validates real-world performance. Dynamic governance—fuelled by weekly KPI reviews, case-based learning, and AI-assisted incident management—drives continuous refinement of data quality, prompts, and model thresholds. Clear ownership by engineering, product, and data governance teams maintains discipline, while user feedback shapes iterative improvements. Expansion proceeds cautiously with strong guardrails, cost controls, and transparent dashboards, ensuring safe, scalable, and enterprise-ready adoption of advanced AI capabilities.

A holistic project roadmap

Illustrative

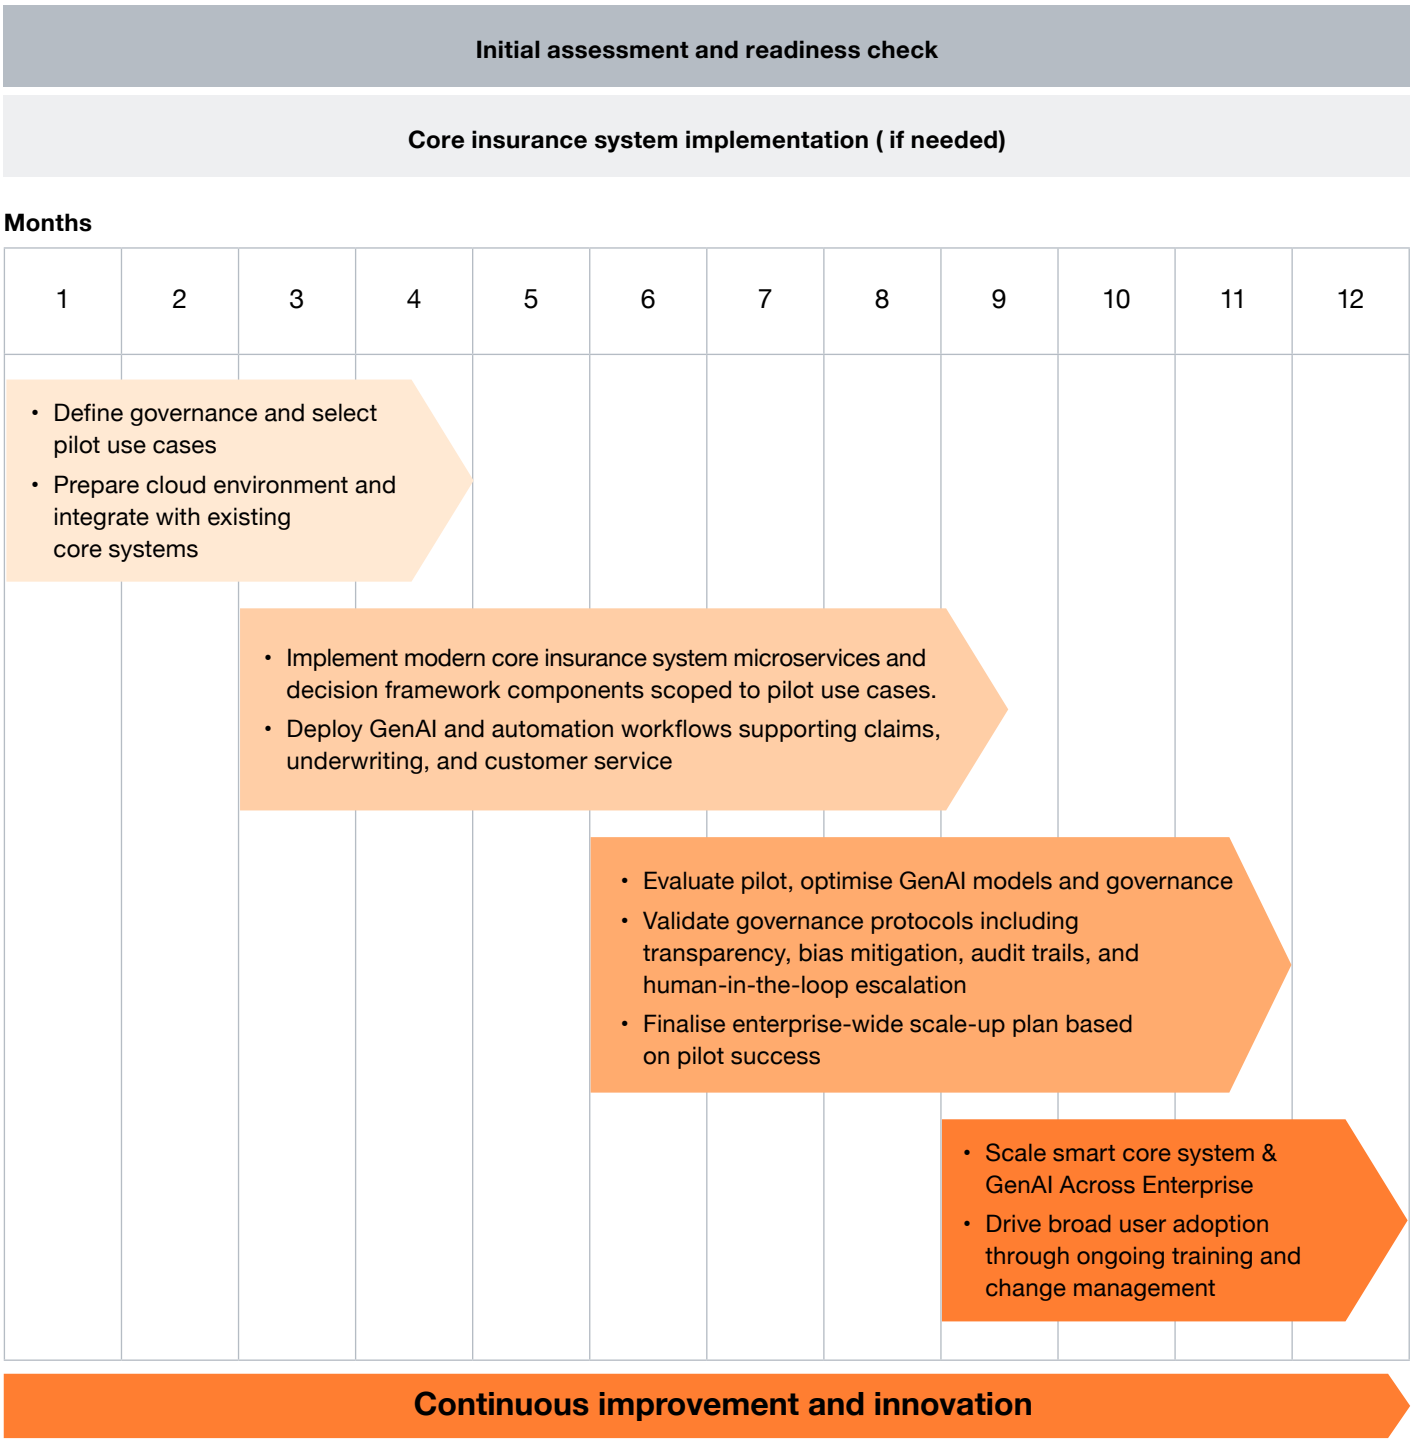

The journey towards a modern, intelligent insurance infrastructure involves measured pilots, collaborative implementation, and sustained governance. A continuous learning loop eventually paves the way for sustainable innovation and resilience.

08

Next is now

Redefining data and decisions towards sustainable outcomes

Insurance is about protecting what matters most, including homes, vehicles, businesses, and people's peace of mind. Behind every policy and claim are real individuals like Polly, Mike, and Sarah. Their everyday experiences shape the industry.

The future of insurance looks promising because it relies on smarter decision-making. But smarter decisions don't just come from faster computers or clever algorithms; they come from connecting trustworthy information and people in ways that work for everyone.

Imagine a world where Polly can update her policy online in minutes, without any frustration or having to repeat herself. Mike, the claims adjuster, can focus on helping customers instead of searching for missing details. Sarah, the underwriter, can confidently assess risks quickly, knowing she is supported by reliable data and clear rules.

This vision is becoming a reality through smart core insurance systems built on modern enterprise platforms. This innovative and dependable technology foundation seamlessly integrates data, processes, and people. With the guidance of a robust decision framework ensuring consistent and fair processes, insurance teams can deliver faster, clearer, and more accurate outcomes. Customers receive answers when they need them, and employees are empowered with the tools to do their best work.

This transformation transcends a mere technical upgrade; it represents a new way to deliver trust, transparency, and speed in an industry that affects everyone's life. Companies that embrace this change will build lasting relationships with their customers while unlocking new levels of efficiency and excellence.

Illustrative

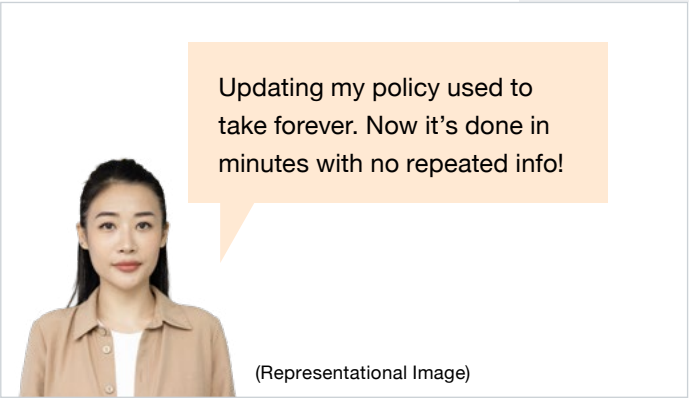

The journey has already begun. The next era of insurance is about blending innovation with empathy, intelligence with clarity, and speed with reliability. By leveraging smart core systems and disciplined decision frameworks, we can create an insurance experience that is smarter, fairer, and truly focused on what matters the most—people.

Successful insurers will be those that thoughtfully invest in these intelligent technologies and frameworks, balancing automation with human insight. This approach enables not only operational excellence but also builds customer trust and drives sustainable growth in an evolving market. PwC is committed to partnering with insurers on this transformative journey to unlock lasting value for businesses and the people they serve.

Decide with confidence: Where decision intelligence meets action

01 Control data at source

Establish one reliable source of truth to ensure consistency and accuracy across all insurance systems and processes.

02 Decide in minutes

Enable faster underwriting and decision-making. This reduces turnaround times to meet customer expectations for quick responses.

03 See risk in real time

Gain immediate visibility into risk factors during quoting, binding, and claims. This helps with proactive risk management.

04 Select the right risk

Use smart triaging to assess and prioritise risks correctly, ensuring focus on high-value and complex cases.

05 Price with precision

Provide data-backed pricing that reflects fair and competitive premiums based on risk profiles.

06 Improve portfolio health

Boost overall portfolio performance by identifying trends, managing exposure, and optimising risk distribution.

07 Embed anywhere

Integrate smoothly across different channels, teams, and partners to deliver consistent decision-making experiences.

08 Scale with confidence

Automate governed decisions at scale while keeping compliance, transparency, and trust in all processes.

Glossary

Term	Definition
Artificial intelligence (AI)	Computer systems that perform tasks typically requiring human intelligence, such as learning, reasoning, and decision-making
Automation	Using technology to perform tasks with minimal human intervention, improving efficiency and consistency
Bias mitigation	Strategies and practices to reduce unfair bias in AI models and data processing
Compliance	Adherence to laws, regulations, and internal policies governing insurance and data usage
Data contract	An agreement outlining the responsibilities, structure, and quality expectations related to specific data products shared across teams
Data mesh	A decentralised data architecture enabling domain teams to manage and share governed data products independently, improving agility and data quality
Data product elements	Components defining a governed dataset including schema, ownership, access rules, quality controls, and metadata used in insurance data platforms
Data stewardship	Responsible management and oversight of an organisation's data assets to ensure quality, privacy, and compliance.
Decision framework	A structured system that applies consistent, traceable decision logic in insurance processes to ensure reliable, fair, and auditable decisions
Explainability	The ability of AI systems to provide understandable reasons behind their decisions to users and auditors
Fallback paths	Predefined procedures taken when AI is uncertain, including human escalation or requesting more information to avoid incorrect decisions
First notice of loss (FNOL)	The initial notification from a policyholder reporting an incident or damage that may lead to a claim
Generative AI (GenAI)	Advanced AI models that generate human-like text or content, assisting in tasks such as summarising policy coverage and drafting communications
General Data Protection Regulation (GDPR)	EU regulation that governs data privacy and protection for individuals within the European Union, ensuring personal data is handled securely and transparently
Guided decision checklist	Dynamic checklists that guide claims handlers and underwriters through policy and claim rules, ensuring compliance and reducing errors

Term	Definition
Health Insurance Portability and Accountability Act (HIPAA)	US regulation protecting sensitive patient health information from being disclosed without the patient's consent or knowledge
Human-in-the-loop	A process where human judgment is included in AI decision-making workflows to ensure oversight and accountability
Key performance indicators (KPI)	Metrics used to measure effectiveness, efficiency, and success of insurance and AI processes
Metadata governance	Management of data about data structures, access rules, and quality to maintain consistent and reliable information across systems
Observability	Monitoring systems that track logs, metrics, and alerts to detect anomalies and ensure system reliability
Payment Card Industry Data Security Standard (PCI DSS)	Standards to secure credit card transactions and protect cardholder data across systems handling payments
Prompt templates	Prewritten AI instructions that specify inputs, reasoning steps, and output formats to ensure consistent and safe AI responses
Role-based access control (RBAC)	Security method that restricts system access to authorised users based on their roles within the organisation
Role-aware service bot	AI assistants tailored to user roles, providing relevant support to claims adjusters, underwriters, or customers.
Smart core system	The core insurance platform that integrates data, technology, and processes to provide a stable, governed base for insurance operations.
Single sign-on (SSO)	Authentication process allowing users to access multiple systems with one set of login credentials
Subrogation	Process where an insurer recovers costs from a third party responsible for an insured loss after compensating the policyholder
Underwriting intake	Collection and verification of information necessary to assess insurance risk and issue policies

About PwC

We help you build trust so you can boldly reinvent

At PwC, we help clients build trust and reinvent so they can turn complexity into competitive advantage. We're a tech-forward, people-empowered network with more than 364,000 people in 136 countries and 137 territories. Across audit and assurance, tax and legal, deals and consulting, we help clients build, accelerate, and sustain momentum. Find out more at www.pwc.com.

PwC refers to the PwC network and/or one or more of its member firms, each of which is a separate legal entity. Please see www.pwc.com/structure for further details.

Contact us

Amit Roy
Partner and Leader—Insurance and Allied Businesses
PwC India
roy.amit@pwc.com

Rajarshi Ghosh
Partner—One Consulting
Data and Analytics, Advisory
PwC India
rajarshi.Ghosh@pwc.com

Authors

Mousumi Sahoo
Priyanka Karri
Arnab Chakraborty

Contributors

Parinita Gupta
Ronit Roy
Samikshya Samantaray
Mustafa Asy
Gitika Karagwal

Editors

Rashi Gupta
Rubina Malhotra

Designer

Harshpal Singh

pwc.in

HS/April 2026 - M&C 52565