November 2026
InfoBrief
Sponsored by NTT DATA

Accelerating Competitive Advantage in Insurance: Private and Sovereign AI

How to Plan, Design, and Implement AI Factories to Meet Sovereign AI Objectives



Inci Kaya
Senior Research Manager, Worldwide Digital Strategies, Insurance, IDC

Table of contents

Click any title to navigate directly to that page.

Across global markets, enterprise adoption of AI is accelerating as it becomes more critical to business success 3

AI-driven transformation in insurance 4

Sovereign AI and private AI: Rewriting how AI is built, governed, and scaled 5

Sovereign AI and private AI: Where the complexity shows up 6

The automation gap must be bridged to enable scaling AI successfully across private and public AI 7

In insurance, carriers redesigning for private and sovereign AI today are creating a lasting advantage 8

De-risking AI adoption in insurance with sovereign AI 9

Engineered AI factory approaches can help simplify, speed up, and secure enterprise AI solutions 10

Private and sovereign AI in action: Strategic industry use cases in insurance 11

Private and sovereign AI in action: Industry use cases in insurance 14

Principles for a carrier-grade sovereign AI capability 16

From pilot to production: Steps to full-stack sovereign AI in insurance 17

Expert partners help accelerate sovereign AI time to value, reduce risk, and boost ROI 18

About the IDC analyst 19

Message from the sponsor 20

Across global markets, enterprise adoption of AI is accelerating as it becomes more critical to business success

AI is driving the **strongest IT spending growth** since 1996.

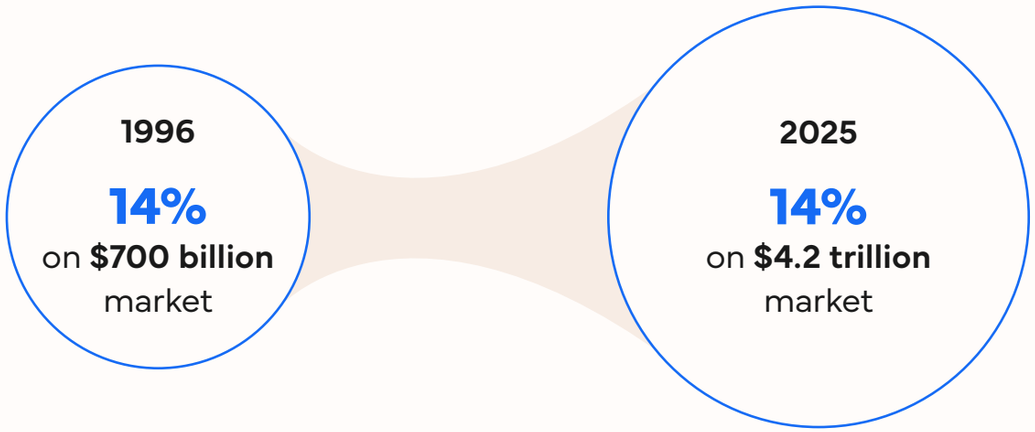


Figure 1

5x growth in **enterprise AI** investment



AI is the most significant IT infrastructure buildout in the history of IT, set to reach $1 trillion by 2029.

In 2026, many enterprises are seeing the impact of AI on their markets and business and are investing in their own AI use cases, as well as private AI infrastructure on which to run them.

Note: Data represents the total worldwide market for IT and enterprise AI spending.
Sources: IDC Worldwide Black Book, March 2026 Forecast; IDC Quarterly AI Infrastructure Tracker, 2026

AI-driven transformation in insurance

Three key drivers are shaping the transformation of the insurance industry:



Enterprise-scale deployment on deck

Carriers are moving from proof-of-concept experimentation to enterprise-scale deployment, bringing governance, cost, and infrastructure decisions to a head simultaneously.



Value-chain remodel

AI is reshaping every part of the insurance value chain (FNOL triage, claims processing, fraud detection, underwriting, actuarial pricing, and compliance reporting), affecting personal lines, commercial lines, specialty/E&S, life, and reinsurance.



AI and the infrastructure revolution

AI is the most significant IT infrastructure buildout since 1996. Enterprise AI investment will reach $1 trillion by 2029, with insurance among the highest-spending verticals in financial services.

AI adoption in insurance is growing rapidly.

2025

63%

Nearly two-thirds of organizations are more likely to adopt sovereign AI due to geopolitical events.

2026

33%

Only a third of insurers use sovereign cloud solutions today.

2026

67%

More than two-thirds of insurers use sovereign AI for compliance with industry regulations.

Private and sovereign AI for insurance: Purpose-built private and sovereign AI infrastructure for training, fine-tuning, and operating domain-specialized AI models keeps policyholder data, claims records, and actuarial inputs under the carrier's control and within the jurisdictions its regulators require.

Source: *IDC Syndicated Survey 2026: Digital Sovereignty Survey*

Sovereign AI and private AI: Rewriting how AI is built, governed, and scaled



Sovereign AI

A country or organization's **capacity to build and run AI on its own terms,** using its own infrastructure, data, models and talent, all anchored in national priorities, laws, and culture



Private AI

AI that runs in an environment **the organization controls** (on-premises and fully air-gapped, in private cloud, or in an accredited sovereign cloud region) alongside its data, with full governance, security, and ownership

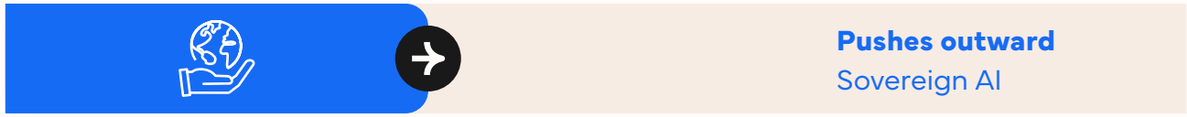


75%

of enterprise AI workloads will run on fit-for-purpose hybrid infrastructure by 2028.

Note: Worldwide market
Source: *IDC FutureScape: Worldwide Digital Infrastructure 2025 Predictions*

Sovereign AI and private AI: Where the complexity shows up



Pushes outward
Sovereign AI

Complexity spans regions, providers, and regulatory regimes.

Pulls inward
Private AI



Complexity stays inside and means securing environments, integrating systems, and keeping internal data flowing cleanly.

Running **AI in either** of these **environments** means **integrating** infrastructure, cloud, data, models, and governance across **hybrid environments:** on-premises, private cloud, sovereign hyperscaler cloud, public cloud, and edge, selected workload by workload, with classification level deciding which are permitted.

Navigating the complexity

The harder organizations push for control, the more interdependent their AI ecosystems become. It is not about owning every piece. Instead, it's about making a multi-provider system run reliably inside defined boundaries. Progress comes from ecosystems with tighter orchestration, clearer accountability, and stronger alignment across partners.

The automation gap must be bridged to scale AI successfully across private and public AI

Agile IT based on a modern digital platform

Inflexible IT based on a legacy architecture

1 **Digital Leaders**
Top 26%
Digitally transformed

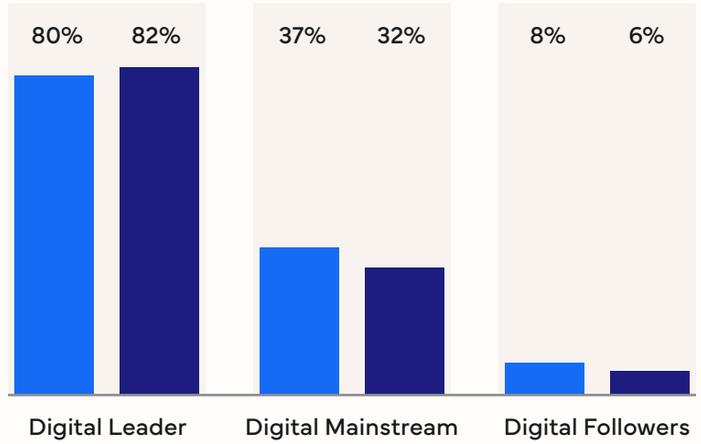
2 **Digital Mainstream**
Middle 49%
Digitally transforming

3 **Digital Followers**
Final 26%
Digitally resisting

Percentage of respondents:

Extensive adoption of automation and orchestration

Extensive use of AI for business use cases



Organizations that have successfully scaled AI use cases have typically extensively deployed:

- Automation and orchestration
- AIOps
- Cloud-native development practices
- Strong security and risk management
- Hybrid and multicloud adoption
- Integrated real-time data

Solutions designed to enable and integrate these capabilities will help accelerate and derisk AI use case deployments.

Digital Leaders invest broadly and deeply across their IT infrastructure, with strategic support and funding; this pays dividends with demonstrably higher business agility and revenue growth. Companies in the Digital Mainstream and Digital Follower segments that are striving to become Digital Leaders should engage with proven transformation partners to assess their current capabilities and plan a journey together to deliver this transformation.

The architecture gap

In insurance, carriers redesigning for private and sovereign AI today are creating a lasting advantage

A clear divide is opening between insurance carriers that have redesigned their AI architecture by incorporating sovereignty, governance, and public and private AI foundations and those that are still running isolated pilots on infrastructure never built to scale. The impact of this gap extends beyond operational aspects: It also encompasses financial and regulatory dimensions.

	Leaders	Followers
AI strategy	Enterprise-wide, targeting full operating model transformation	Select functions or no enterprise AI strategy
Data governance	Automated, full traceability and real-time controls	Fragmented, inconsistent, still being built
Infrastructure	Purpose-built, public and sovereign foundations in place	Reactive deployment, minimal scalability planning
Risk management	Dynamic, real-time risk management	Manual reviews, inconsistent human oversight
AI road map	Agentic workflows, cross-functional orchestration	Isolated pilots, short-term use case focus

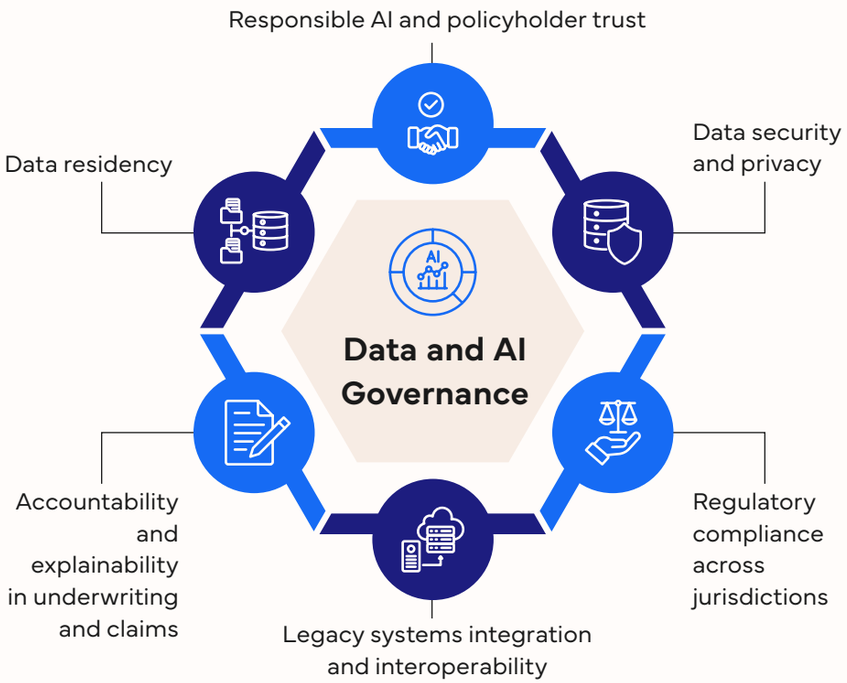
The cost of waiting

The compliance landscape and AI use case expectations are constantly evolving. Organizations running AI on ungoverned architectures face a growing challenge: Each new regulation, client demand, or market factor increases the distance between their current state and where they need to be in terms of technology, operations, or institutional readiness.

Sources: IDC *AI MaturityScape Benchmark 2026,* April 2026 (n=234); IDC *AI Infrastructure Evolution Survey,* March 2025 (n=75)

De-risking AI adoption in insurance with sovereign AI

Six AI governance dimensions



Insurance de-risking AI

Sound AI governance that ensures fairness, explainability, and accountability is the foundation for safe AI at scale. **Insurers must de-risk AI adoption by addressing:**

- **AI control:** Domain expertise to prevent bias and full model control over high-risk underwriting and claims decisions
- **Governance:** Revised processes and differentiated governance frameworks across claims, underwriting, and actuarial functions, including model approval, monitoring and observability, human oversight, incident management, auditability, and regulatory evidence
- **Compliance:** A cross-regulation, cross-jurisdiction challenge spanning the EU AI Act, DORA, GDPR, and national or state insurance overlays, including BaFin, EIOPA, VAG, IVASS, PRA/FCA, NAIC model laws, and relevant local market requirements
- **Controlled environments:** A shift toward private and sovereign environments, alongside public cloud where requirements allow, to retain control over policyholder data and AI model behavior
- **Security concerns:** Growing insurer focus on data security and privacy as AI expands across claims, underwriting, and fraud detection
- **Sovereignty rise:** Rising insurer interest in sovereign infrastructure, driven by geopolitical tension and tightening regulatory pressure

Insurers must manage these requirements across jurisdictions and lines of business, which affects the overall architectural choices supporting AI deployments. Policyholder data-residency requirements, while not unique to AI, will directly shape how the AI factory architecture is engineered.

Private and sovereign AI provide the technological foundation to support compliance readiness and help address relevant regulatory requirements, retain model control, and deliver competitive value securely and at scale across jurisdictions and lines of business.

Engineered AI factory approaches can help simplify, speed up, and secure enterprise AI solutions

AI factories are:

- Industrialized, scalable digital infrastructure platforms
- Purpose-built to automate and operationalize the development, deployment, and management of AI at enterprise or national scale across AI-related hardware, software, and services
- Vertically integrated platforms with specialized AI accelerators, fabrics, orchestration software, AI/ML frameworks, and cloud/hybrid infrastructure
- Able to deploy a combination of commercial or open AI models from large to small to support multi-modal solutions, including general-purpose LLMs, optimized SLMs, and diffusion models
- Able to streamline the transformation of enterprise data into actionable intelligence

Benefits which you can expect:



Accelerate time to value



Cut costs



Reduce risk

Expert partners with end-to-end engineered AI factory platforms that are right-sized for your private AI workload needs will be critical to meeting business value outcomes and overall objectives.

Private and sovereign AI in action: Strategic industry use cases in insurance

Use case 1:

Underwriting and actuarial services

Private and sovereign AI enable insurers to train models on proprietary loss and exposure data. It delivers precision risk pricing and helps address relevant EU AI Act requirements for automated decisions.

- **Personal lines:** Individual-level actuarial modeling for tailored pricing and risk segmentation
- **Commercial, specialty, E&S:** Early detection of emerging risk concentrations before they escalate into catastrophic exposures
- **Competitive moat:** Proprietary models that remain confidential and cannot be replicated by competitors
- **Regulatory alignment:** Full model control, explainability, and auditability by design

By leveraging an AI factory, insurers can orchestrate agentic workflows and dynamic risk management while continuously retraining models on live portfolio data for near-real-time risk reevaluation, preserving confidentiality and regulatory compliance. This creates a competitive moat, as proprietary models cannot be replicated by competitors and remain under the carrier's control.

Private and sovereign AI in action: Strategic industry use cases in insurance

Use case 2:

Claims processing and FNOL

Claims handling is the critical "moment of truth" for any insurer. Sovereign AI transforms the life cycle from FNOL intake through investigation, reserving, and payment, using natural language processing (NLP) to classify and route claims, surfacing subrogation recoveries that manual review misses, and compressing cycle times to lift policyholder satisfaction, all within the controls the carrier defines.

- **NLP-driven automation:** Classifies and routes claims, surfaces subrogation opportunities, and compresses cycle times
- **Policyholder experience:** Faster, more accurate claims handling, boosting satisfaction and trust
- **Compliance:** Supports GDPR compliance readiness and relevant sector-specific data residency requirements; helps address relevant DORA and EU AI Act requirements for automated decisioning
- **Continuous improvement:** Proprietary claims data that retrains models for ongoing accuracy and fraud reduction

During catastrophic loss events, when claim volumes spike and AI-assisted triage matters most, a sovereign architecture is designed so claims processing does not depend on any single external cloud provider or API being available, sustaining consistent policyholder response even under peak-demand conditions.

Private and sovereign AI in action: Strategic industry use cases in insurance

Use case 3:

Fraud detection and investigation

Sovereign AI is business critical for fraud detection; models trained on carrier-specific fraud patterns outperform generic systems and keep sensitive data in-house.

- **In-perimeter training:** Builds unique pattern libraries, controlling exposure to public-API inference risks
- **Real-time defense:** Agentic AI enables continuous threat detection and rapid response
- **Regulatory compliance:** The architecture helps address relevant DORA and EU AI Act requirements for explainability, auditability, and resilience
- **Operational resilience:** Reduces false positives and supports proactive risk management

In fraud detection, the sovereign imperative is most acute. Models trained on a carrier's own historical fraud patterns consistently outperform rules-based and generic systems. Sending raw claims data to uncontrolled external APIs for fraud scoring exposes the very data that must be protected. By keeping training and inference on infrastructure the carrier controls, on-premises or in a sovereign cloud region, insurers build carrier-specific pattern libraries that cannot be reverse-engineered or exposed to public-API inference risk. IDC research highlights that agentic AI trained on proprietary telemetry enables continuous, real-time threat detection, reducing false alerts by up to 60% and supporting proactive risk management.

This sovereign approach also alignswith DORA and the EU AI Act, ensuring that fraud detection systems are explainable, auditable, and resilient to adversarial threats.

Private and sovereign AI in action: Industry use cases in insurance

Where sovereign AI delivers the greatest impact across the value chain

Use case 4:
Compliance, audit, and regulatory reporting

Sovereign AI architectures turn compliance from a cost center into a strategic capability, automating evidence packs, audit trails, and regulatory reporting across jurisdictions.

- **Multijurisdictional alignment:** This navigates EU AI Act, DORA, GDPR, and national overlays (BaFin, IVASS, PRA/FCA, NAIC, state commissioners)
- **Design constraint:** Compliance dictates where models run, how data moves, and any performance trade-offs
- **Automated governance:** Unified, real-time controls ensure operational resilience and regulatory readiness
- **Competitive advantage:** Organizations with automated compliance outpace those layering it on later

For carriers operating across jurisdictions, compliance is a tangle of overlapping, if not outright conflicting, requirements:

- **The EU AI Act** may classify certain automated insurance decisioning use cases as high risk, depending on their purpose and application, requiring applicable auditability, human oversight, and technical documentation
- **DORA** mandates operational-resilience testing, third-party ICT risk management, and supervisory evidence packs
- **GDPR or data protection by design** restricts policyholder personally identifiable information (PII) from transiting external AI APIs without explicit consent frameworks
- **National and regional overlays:** BaFin, EIOPA, VAG, IVASS, PRA/FCA, NAIC model laws, state insurance commissioners, and relevant local market requirements all add requirements

Across a multi-jurisdiction footprint, this matrix becomes unmanageable without AI assistance, turning compliance interpretation into a capability rather than a cost center. This is not only a compliance matter but also a design constraint on where models can run, how data is allowed to move, and the performance trade-offs that follow.

Private and sovereign AI in action: Industry use cases in insurance

Where sovereign AI delivers the greatest impact across the value chain

Use case 5:
Bancassurance and embedded insurance data governance

Where insurance and banking products coexist, sovereign AI enforces strict data separation and access controls, supporting regulatory compliance readiness for cross-sell AI initiatives.

- **Logically separated environments:** Attribute-based access controls and clean audit trails demonstrate compliance
- **Differentiated governance:** AI governance standards tailored by infrastructure jurisdiction are now the norm
- **Ecosystem integration:** AI factories enable secure, compliant, and auditable cross-sell and embedded insurance offerings
- **Regulatory assurance:** This helps address relevant requirements for data-use restrictions and supports ongoing regulatory change

Where insurance and banking products are offered under one institutional roof, regulators require strict separation of insurance and banking data. Sovereign AI architectures enforce logically separated model environments and attribute-based access controls, producing clean audit trails that demonstrate compliance with data-use restrictions.

AI factories support this by enabling end-to-end ecosystem integration, systems interoperability, and regulatory-aligned reference architectures. This ensures that cross-sell AI initiatives in bancassurance and embedded insurance remain compliant, secure, and auditable, even as they leverage advanced AI for hyperpersonalized product offerings.

Principles for a carrier-grade sovereign AI capability

Data stays under carrier control	Training data, model inference, and output generation run in environments the carrier controls, on-premises, in private cloud, or in a sovereign cloud region. Policyholder PII, claims data, and actuarial inputs reach external AI services only where the carrier has decided they can, under its own controls. This embeds data-residency and data-control requirements into the architecture, complemented by appropriate governance, contractual, and third-party risk controls.
Domain-specialized models over generic LLMs	Models fine-tuned on insurance-specific documents, whether ACORD forms, loss-run reports, or similar sources, understand distinctions that general-purpose LLMs miss. This translates to higher accuracy at lower inference cost for carriers while simultaneously protecting that proprietary training data as institutional IP.
Explainability by design for high-risk decisions	The EU AI Act's Annex III identifies AI systems to be used for risk assessments and pricing as high risk because of the profound impact these factors can have on people. A sovereign architecture can generate explainability artifacts and provide feature-level rationale, data lineage, and model-version records, making audit a continuous background process rather than a one-off crisis response.
Operational resilience under DORA	Critical inference runs on infrastructure the carrier controls, so no single cloud-provider or API outage can stop claims triage. During a catastrophic loss event, just when claims volumes spike and AI-assisted triage matters most, the carrier is not dependent on one external provider's availability or SLA, which is the ICT concentration risk DORA asks it to manage.
Multi-jurisdiction compliance automation	For MGAs, reinsurers, and multinational commercial carriers, programmatic interpretation of regulatory requirements can automate form selection, data-residency routing, and supervisory reporting at the transaction level, converting a staff-intensive process into an auditable, AI-governed workflow.
Orchestration across the hybrid estate	Full-stack orchestration, spanning on-premises, private cloud, sovereign hyperscaler cloud, and public cloud, lets carriers place each workload where its risk, jurisdiction, and cost dictate, while still integrating with legacy policy and claims systems, rather than forcing an all-or-nothing infrastructure choice.

A carrier-grade sovereign AI capability rests on a set of architecture principles tuned to insurance data, regulation, and risk culture.

From pilot to production: Steps to full-stack sovereign AI in insurance

1 Assess and map AI exposure across the value chain

Map every AI workload, from FNOL triage, claims scoring, underwriting models, and actuarial engines to fraud detection against data sensitivity, regulatory jurisdiction, and line of business. P&C, life and annuities, reinsurance, and specialty lines carry different risk profiles. Build an architecture blueprint per line, not a single compliance checklist for all.

2 Build data governance and sovereign AI foundations first

Carriers cite integration with legacy policy and claims systems as their number-one deployment challenge because they retrofit governance onto infrastructure not designed for AI. Data sovereignty, in-perimeter training pipelines, policyholder PII controls, and ongoing regulatory and compliance or assurance requirements must be in place before AI moves from the actuarial sandbox into production underwriting or live claims.

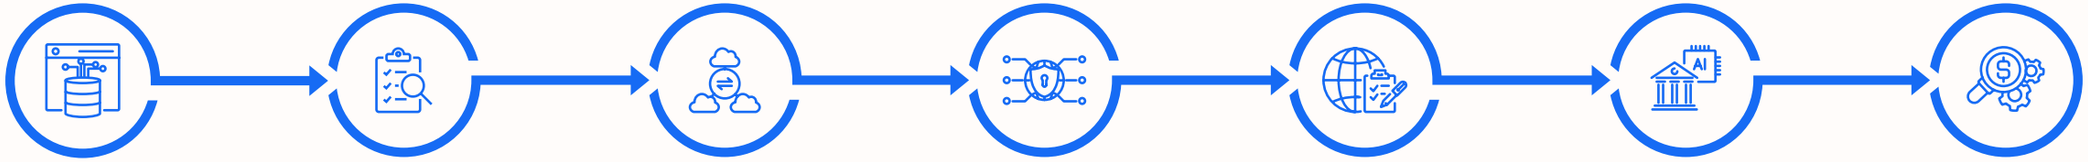
3 Industrialize through the insurance AI factory

Move from bespoke carrier POCs to a standardized, repeatable platform: shared MLOps pipelines, domain-specialized models fine-tuned on ACORD forms and loss-run data, unified claims governance, and hybrid orchestration across on-premises, private cloud, sovereign hyperscaler cloud, public cloud, and edge.

4 Apply sovereignty where risk demands it; scale progressively

Reinsurers, MGAs, and multinational commercial carriers face high jurisdictional complexity. Life and annuities carriers handle the most sensitive long-duration policyholder data. P&C claims face the tightest real-time auditability demands. Without a risk-tiered deployment model, sovereign AI investments are made blindly, overengineering low-risk workloads and underprotecting high-risk ones.

Sovereign AI for insurance spans infrastructure, data engineering, domain AI models, compliance, and change management, making it impossible to build alone. Choose a partner that combines TPA operational depth, regulatory intelligence across jurisdictions, validated AI factory blueprints, and multiprovider orchestration. Look at their capabilities in relation to:



Full-stack coverage
From datacenters to AI services

Regulatory knowledge
DORA, EU AI Act, GDPR, and insurance regulations across jurisdictions, including third-party risk, audit or access rights, resilience testing, subcontracting, exit, and supervisory evidence

Multi-LLM and multicloud
Avoiding vendor lock-in and enabling geographical resilience

Security and safety
Operational resilience, cybersecurity

Insurance domain expertise
ACORD, Axco, and loss-run models with TPA operational depth

AI model governance
Transparency by design

TCO and ROI
Frameworks for optimized deployment cost

Expert partners help accelerate sovereign AI time to value, reduce risk, and boost ROI

Consulting, integration, and operations services for private and sovereign AI

- Workload assessment and placement across on-premises, private, sovereign, and public cloud
- Vendor-agnostic, turnkey, prevalidated infrastructure deployment
- Integration and operations across hyperscaler, private AI, edge, and legacy estates
- AI strategy and roadmap
- Feasibility and readiness assessments
- Data governance and strategy
- Use case prioritization
- Data pipeline engineering
- MLOps and model management
- Dealing with skills shortages

Enterprise benefits

- Accelerated time to deploy with validated blueprints, reference architectures, repeatable methodologies, and technical expertise
- Resource optimization and utilization for greater TCO/ROI across private, public, and on-premises locations (FinOps best practices)
- Reduced operational risk through improved configuration, expert resources, and knowledge transfer sessions
- Improved security and compliance through secure firewalls and deep understanding of regulated industries

Partners can help pinpoint the areas of greatest business impact in your organization and develop roadmaps for responsible, enterprise-wide AI transformation, addressing every dimension of privacy and sovereignty.

About the IDC analyst



Inci Kaya

Senior Research Manager, Worldwide Digital Strategies, Insurance, IDC

Inci Kaya is Senior Research Manager for IDC Financial Insights, responsible for digital strategies for insurance worldwide. Mrs. Kaya's core research coverage includes the evolution of AI and Agentic AI, fraud prevention and cybersecurity, claims processing themes, risk and compliance issues in insurance. Based on her background in technology and financial services space, Mrs. Kaya's research also includes elements of digital payments, customer experience, cloud services and data analytics.

More about Inci Kaya →

Message from the sponsor



The value of AI in insurance is inseparable from the trust placed in it. As AI takes on more responsibility across pricing, underwriting, and claims, the ability to govern, audit, and explain every model decision matters as much as the outcomes it delivers.

Sovereignty is a set of requirements, not an infrastructure model. Policyholder data, actuarial models, and claims decisions each carry their own regulatory, security, and residency obligations, and no single environment serves them all. NTT DATA helps insurers decide what runs where: on-premises, in private cloud, in a sovereign hyperscaler region, or in public cloud, workload by workload, under the jurisdiction and controls each demands. We then integrate and operate the whole environment, across hyperscalers, private AI, edge, and legacy policy and claims systems, as one governed estate.

Our AI Factory is the operating model behind this: validated reference architectures, insurance-grade ontologies, and embedded governance that protect proprietary actuarial models, deliver explainable AI in high-risk functions, and monitor continuously for bias and drift. With EIOPA, the EU AI Act, and NAIC frameworks raising expectations, auditability, human oversight, and data residency are designed in from the outset, so insurers can move from fragmented pilots to production AI that regulators, boards, and policyholders can trust.

Learn how NTT DATA helps insurers operationalize private and sovereign AI securely, responsibly, and at scale, and explore our private and sovereign AI solutions.

IDC Custom Solutions

IDC Custom Solutions produced this publication.
The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis that IDC independently conducted and published, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies.

This IDC material is licensed for external use and in no way does the use or publication of IDC research indicate IDC's endorsement of the sponsor's or licensee's products or strategies.

IDC.com | IDC on LinkedIn | IDC Blog

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries. IDC's analysis and insight helps IT professionals, business executives, and the investment community to make fact-based technology decisions and to achieve their key business objectives.

©2026 IDC. Reproduction is forbidden unless authorized. All rights reserved. CCPA