NTT DATA

Closing the execution gap in insurance

Activating systems of execution with agentic AI

Driving intelligent change for insurance

From insight to execution

This Everest Group paper, supported by NTT DATA.

Insurers are making significant progress in generating data and insights, yet many find it difficult to translate these into consistent, scaled outcomes across the enterprise.

At NTT DATA, we believe the next phase of transformation will be defined by the ability to operationalize decision-making at scale and orchestrate intelligent workflows across core processes. This paper explores how systems of execution, powered by agentic AI, can enable more effective, real-time execution.

Driving intelligent change for insurance

Partnering with NTT DATA enables insurers to operationalize AI, orchestrate intelligent workflows and move from insight to execution at scale.

Everest Group®

April 2026

Systems of Execution (SoE) in Insurance: Activating Agentic AI to Close the Execution Gap

Contents

03 Introduction

04 Insurance does not have a data problem – it has an execution gap

05 Agentic AI in the context of SoE

06 The execution architecture for agentic insurance

08 An SoE-led reinvention model

09 SoE in action across insurance use cases

14 An SoE adoption roadmap

16 Constraints that will slow autonomy

19 Conclusion

Introduction

The insurance industry is gradually transitioning from static processing to autonomous orchestration. For decades, carriers invested heavily in Systems of Record (SoR) to ensure data integrity across policy, underwriting claims, and billing environments. While these systems established a single version of truth, they were never designed to activate intelligence in real time. With carriers adopting advanced analytics and generative AI, a new constraint has emerged: the inability to translate insights into governed, scalable execution within core workflows.

Systems of Execution (SoE) address this constraint by acting as an orchestration layer that connects intelligence to action, enabling AI agents to operate safely across legacy cores, third-party ecosystems, and workflow engines. By embedding governance, authority thresholds, Human-in-the-Loop (HITL) controls, and observability into execution paths, SoE enable carriers to move from insight at rest to execution in motion.

In this Viewpoint, we examine:

- The structural execution gap limiting enterprise AI value realization
- Agentic AI's role in transforming insurance workflows
- SoE's architectural foundations
- Segment-specific implications across Property and Casualty (P&C) and Life and Annuities (L&A) insurance
- A pragmatic roadmap for scaling governed autonomy

This report is intended for insurance enterprise leaders, technology executives, and operations stakeholders seeking to responsibly industrialize agentic AI. It outlines how carriers can move beyond pilot experimentation and build the execution architecture required to unlock sustainable enterprise value.

Insurance does not have a data problem – it has an execution gap

The fundamental limitation in insurance is not the lack of data, but the inability to activate that data for decision-making. Traditional core systems operate on linear, deterministic logic that is incompatible with AI-driven work's non-linear and probabilistic nature. The insurance industry has historically relied on point solutions to address business challenges, but these have largely failed to address the end-to-end execution life cycle. For example, while predictive models can flag a potentially fraudulent claim, the investigation – gathering evidence, cross-referencing databases, and initiating a referral – remains a manual and fragmented process. This is a structural barrier where insights are generated in silos, requiring human intervention to manually bridge those insights back into systems for appropriate action execution. This gap between insight and action keeps widening as risks become more complex and data-intensive.

The execution gap manifests differently across insurance segments. In P&C insurance, variability stems from loss events, vendor coordination, catastrophe exposure validation, and jurisdictional claims regulations. In L&A insurance, friction is concentrated in long-cycle underwriting, medical evidence ordering, suitability validation, beneficiary changes, and policy servicing events. While the architectural constraint is common, the execution patterns and control intensity differ materially by segment.

Closing this gap requires embedding intelligence directly into operational workflows, not layering additional tools on top of fragmented workflows. This is where agentic AI emerges as the inflection point.

Agentic AI as the inflection point

Agentic AI unlocks the possibility of autonomous work by shifting AI's role from an assistant to an executor. While generative AI excels at completing manual tasks, agentic AI has the potential to take the next step toward execution. However, for these agents to move from sandbox experiments to enterprise reality, carriers must adopt an architecture that provides agents with the necessary permissions, context, and guardrails to act safely.

That enabling execution architecture is SoE, allowing carriers to unlock AI currently trapped in workbenches and dashboards, and move from insight at rest to execution in motion.

Exhibit 1 illustrates SoE's transformation potential, highlighting the shift from siloed, legacy-constrained operations to integrated, scalable, and execution-driven enterprise models.

Exhibit 1: SoE's transformation potential

Source: Everest Group (2026)

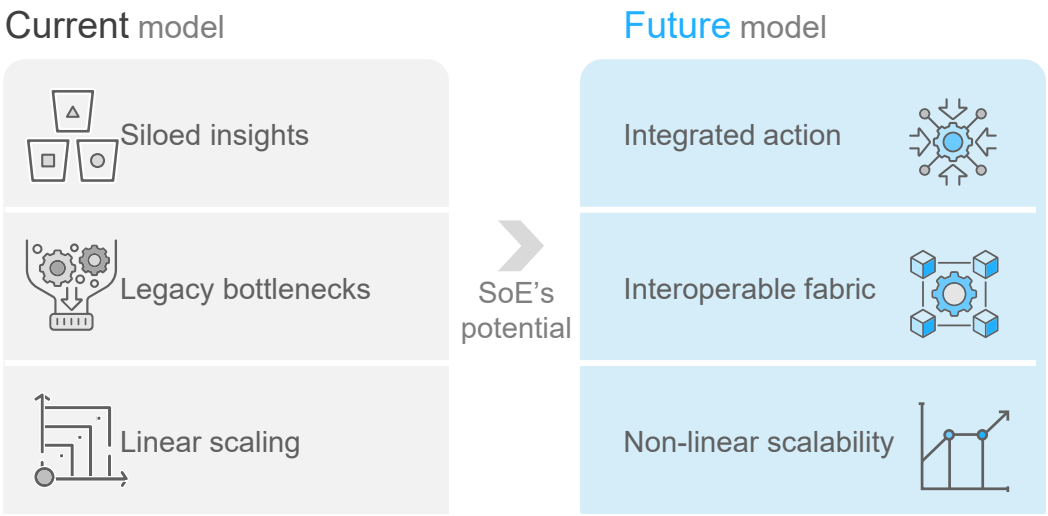

To understand how this shift materializes operationally, it is important to examine what agentic AI means within insurance workflows and how these agents interact with enterprise systems.

Agentic AI in the context of SoE

In operational terms, agentic AI refers to autonomous (or semi-autonomous) systems that can plan, decide, and act across tools and workflows to achieve a defined business outcome.

An AI insurance agent is not just a model producing an answer, but a digital worker that can:

- **Plan:** break an outcome (for example, make submission quote-ready or progress claim to settlement) into a sequence of steps.
- **Decide:** select the right path based on policy rules, case context, authority limits, and confidence thresholds.
- **Act:** invoke tools (such as workflow engines, provider platforms, and document services) to execute steps and move the case forward.
- **Collaborate:** hand-off to humans at defined control points, capture evidence, and resume execution after approvals.
- **Adapt:** adjust the plan based on intermediate outcomes (such as missing evidence, exceptions, fraud triggers, and authority breaches)

This plan-decide-act loop, augmented with collaboration and continuous adaptation, also distinguishes agentic AI from copilots, facilitating the move from value creation in isolation to embedded intelligence connected to execution. This distinction is particularly important across insurance product segments. In P&C insurance, agentic systems can accelerate workflows, including submission intake, risk enrichment, catastrophe response, and vendor coordination in claims. In L&A, the value often lies in orchestrating documentation-heavy journeys, such as underwriting evidence collection, policy servicing, and beneficiary claims processing, where decision accuracy and regulatory defensibility are critical.

Thus, agentic AI's impact can be framed across three dimensions:

- **Productivity:** capacity reallocation enabling skilled professionals to spend more time on judgment, negotiation, exceptions, and customer empathy over administrative execution
- **Explainability:** audit-ready governance in a regulated environment where better decisions are explainable, reviewable, and policy-bound
- **Agility:** compressed cycle time when agents reduce the orchestration tax (the hidden latency between steps)

However, these agents' effectiveness ultimately depends on the environment in which they operate. Without the right execution architecture, even the most advanced agents remain confined to advisory roles.

The execution architecture for agentic insurance

When intelligence can plan and act, the focus shifts from generating the right insights to executing them reliably, safely, and at enterprise scale across the carrier's operating environment.

SoE is the execution backbone that enables agentic AI at scale. It provides orchestration, tool governance, context, controls, observability, and security in a unified architecture.

SoE is best understood through capabilities that make agentic execution possible in insurance. A critical but often overlooked capability in this stack is insurance ontology, a structured representation of risks, policies, exposures, claims events, customers, agents, and loss histories, among other things. Ontologies organize how these entities relate to one another, allowing AI agents to interpret context consistently across underwriting, claims, and servicing workflows.

By embedding domain semantics into the execution layer, carriers can enable agents to reason over insurance-specific concepts rather than generic data objects, improving decision accuracy and workflow orchestration. When stitched together, these capabilities form a reference stack. This stack does not replace SoR; it wraps them in an execution layer that is workflow native, policy-aware, and audit-ready.

Exhibit 2 illustrates six key SoE elements that enable enterprise-grade autonomy.

Exhibit 2: Six SoE elements that enable enterprise-grade autonomy

Source: Everest Group (2026)

Agent runtime and tool layer

Bridges autonomous action through connectors with deterministic logic and strict authority limits to ensure version controlled and compliant executions

Observability and value monitoring

Presents a transparent storyboard of the agentic workforce's operational impact and efficiency by tracking multiple KPIs and identifying model drifts or system failures

Security and privacy guardrails

Ensures cross-border data sovereignty and continuous compliance by modulating access control for agents through strict frameworks

Decision to orchestrate

Decomposes high-level tasks into autonomous chains, maintaining state and synchronizing specialized agents in real time by coordinating complex and multi-agent workflows

Context fabric

Translates raw data into actionable intent, providing agents with the deep situational awareness required for precise decision-making by acting as system's semantic nervous system

Controls by design

Ensures sensitive financial tasks undergo mandatory technical approvals to meet transparency standards by embedding regulatory compliance and audit trails in the workflow architecture

The execution architecture must also account for structural differences across insurance segments. In P&C insurance, workflows are often event-driven and coordination-heavy, requiring orchestration across brokers, vendors, adjusters, and third-party data providers. In contrast, L&A insurance execution is typically documentation-intensive and compliance-bound, involving medical evidence collection, suitability validation, and beneficiary servicing processes. Consequently, the SoE stack must support both high-frequency coordination workflows and evidence-heavy decision workflows while maintaining consistent governance and auditability.

While architecture enables execution, realizing enterprise value requires more than technology design. Carriers must also rethink how these capabilities are delivered and embedded into operational workflows.

An SoE-led reinvention model

AI programs stall, not because the models fail, but because organizations try to scale agentic AI the way they scaled earlier generations of digital initiatives: by layering tools on top of fragmented workflows. This primitive approach creates exactly the conditions frontline teams complain about – too many tools, systems, and manual handoffs. What drives meaningful impact is orchestration, embedding intelligence directly into workflows so decisions can become actions. This is why a combined platform and services model is foundational to implement agentic AI in insurance.

The combined model helps segregate what should be standardized and reused from what must be tailored to the carrier's operating context across two layers:

- **The platform layer:** productizes the engine room (integration, controls, and monitoring) so it becomes reusable across workflows and lines of business
- **The services layer:** embeds forward-deployed squads into underwriting, claims, and servicing to reshape workflows, tune human checkpoints, and operationalize governance

The strategic benefit is clear as carriers stay focused on business outcomes, while the heavy lifting needed to make autonomy safe and scalable is standardized.

The pillars of a reusable SoE platform

A reusable SoE platform should be built on the following pillars:

- **Reference architecture:** a consistent blueprint for how orchestration, agent runtime, context fabric, controls, observability, and security fit together
- **Insurance domain ontology:** a standardized semantic layer codifying relationships between key insurance entities, such as risk, product, exposure, customer, agent, and loss, enabling reusable agent patterns and consistent interpretation across workflows
- **Reusable agent patterns or skills library:** insurance-native patterns that repeat across value streams, such as intake/structuring, enrichment/verification, triage/routing, evidence packaging, and controlled action initiation
- **Control blueprints:** reusable governance templates that codify authority thresholds, approval gates, duty segregation, evidence capture, and audit trails
- **Integration adapters:** standard connectors for core platforms and key ecosystems to reduce point-to-point integration sprawl
- **Observability templates:** pre-defined monitoring and evaluation patterns that show what agents did, where humans intervened, whether behavior is drifting, and what KPI impact is being realized

Capabilities enabling the services layer

Even the best platform will fail if it assumes insurance execution is “configure and go”. Real transformation requires workflow reinvention in the operating environment. This makes the resurgence of Forward-deployed Engineering (FDE) patterns inevitable. Embedded squads work directly with underwriting, claims, and servicing teams to:

- Analyze real workflows (including exceptions and handoffs)
- Embed guardrails and control gates into execution paths
- Integrate legacy cores and partner tools pragmatically
- Tune HITL checkpoints
- Set up governance routines and evidence capture aligned to risk postures
- Instrument value monitoring so outcomes are tracked from day one

Once the execution platform and delivery model are established, the next logical question becomes where SoE will create the earliest and most visible operational impact across the insurance value chain.

The best starting points for SoE are not the most ambitious use cases, but those where execution pathways, control guardrails, and value attribution already exist, making impact both achievable and measurable.

SoE in action across insurance use cases

SoE’s impact is not evenly distributed across the insurance value chain. The quantum would depend on the intensity of execution tax (including handoffs, delays, rework, and inconsistent controls), which, in turn, depends on a combination of three factors – structure of work, cross-system engagement, and degree of control. The impact is highest where workflows are unstructured, coordination-heavy, and control-intensive, because that is where the execution tax is most expensive.

Exhibit 3 illustrates how SoE make tangible impact across the insurance value chain.

Exhibit 3: SoE's execution impact across value chain stages

Source: Everest Group (2026)

Value chain stage	Current process	Execution impact using SoE
Distribution and intake	Unstructured inbound emails, attachments, and calls drive late triage while experts spend time on incomplete/non-viable work	**Unstructured intake becomes executable:** Agents structure requests, detect missing information early, enrich context, and route work based on appetite/priority with traceability
Underwriting and new business	Quote cycles are dominated by data assembly, enrichment, and referral coordination	**Quote-ready by design:** SoE orchestrate intake, enrichment, validation, and referral packaging so underwriters receive decision-ready cases
Policy servicing and billing	Route and wait persists because resolution requires cross-system updates under compliance constraints	**Resolve, not route:** Agents interpret requests, validate eligibility, execute changes under approval policies, and generate compliant documentation
Claims and benefits	Queue-driven progress, sequential evidence gathering, late vendor actions, and downstream controls result in cycle time and leakage	**Event-driven orchestration:** SoE launch parallel workstreams (coverage, severity, fraud signals, vendors, and evidence) and enforce payment controls in the execution path
Fraud, Special Investigations Unit (SIU), and recoveries	Escalation and recovery actions happen late because triggers and evidence packaging are manual	**Earlier, disciplined escalation:** Agents continuously package evidence, trigger referrals based on policy, and orchestrate recovery steps with audit trails

Where it hits first: Underwriting and claims, as these two stages combine high volume, high friction, material outcomes, and repeatable execution patterns that can be standardized into reusable modules.

What is different with SoE: Earlier, generative AI improved outputs (summaries and drafts), but SoE enables embedded, autonomous insights that can move workflows – handling unstructured content, pulling third-party context, improving data readiness, and initiating next-best actions within underwriting/claims systems under controls.

The selection logic: considerations before initiating SoE

A typical way in which organizations fail is when they select AI use cases for on-paper extraordinary capabilities rather than execution readiness and get stuck in pilot loops. The right approach would be to treat SoE as an operating model capability and start only where guardrails, tool access, context, and measurement can be engineered from day one.

A multi-lens evaluation is critical to ensure all dimensions of execution readiness are adequately addressed. Exhibit 4 presents a structured lens to assess execution readiness before initiating SoE.

Exhibit 4: Considerations for organizations across categories before initiating SoE

Source: Everest Group (2026)

Outcome clarity: value defined at the workflow level	What metric will move (such as cycle time, leakage, hit ratio, cost to serve, customer effort, and compliance quality) and where in the workflow? What is the execution bottleneck (such as handoffs, rekeying, evidence chasing, queue latency, and inconsistent approvals)? What does complete look like operationally (quote-ready submission, settlement-ready package, or endorsement executed with audit trail)?
Guardrails and accountability	Can we define stop conditions and guardrails (thresholds, exception rules, or confidence triggers)? Where is HITL mandatory vs. optional across authority decisions, adverse outcomes, disputes, and high severity? Can we enforce segregation of duties (recommend vs. approve vs. execute)?
Execution feasibility	Are there tool pathways for action (such as APIs/connectors or controlled automation fallback where needed)? Can we externalize the workflow from core constraints without destabilizing production?
Context readiness	What is the Minimum Viable Context (MVC) required to act safely? Can we preserve provenance/recency, such as what came from where, when, and which version?
Value mapping	What is the cost to run in terms of HITL oversight load, governance reviews, monitoring/evaluation, integration maintenance, and third-party data costs? Can value be measured at the workflow level instead of the model level with attribution of KPIs and traceability per journey?

The claims function meets the SoE readiness criteria, with clearly measurable outcomes, well-established controls, defined execution pathways, and value attributable at the workflow level, making it an ideal proving ground before scaling across other insurance workflows. We take a closer look at SoE's impact in claims next.

Claims process execution: improving cycle time and control through orchestration

SoE's advantage becomes quickly visible in claims because orchestration quality directly shapes cycle time and leakage. The goal should not be faster task execution but parallelized and event-driven execution with controls embedded in the path.

Claims orchestration also varies materially by segment. In P&C insurance, execution complexity arises from parallel workstreams, such as coverage validation, vendor coordination, severity assessment, subrogation, and fraud screening, often across multiple external partners. In L&A insurance, claims are less frequent but highly sensitive, requiring beneficiary validation, death certificate verification, ownership checks, policy status confirmation, and regulatory reporting. While P&C insurance providers emphasize cycle compression and leakage control, L&A providers prioritize documentation accuracy, beneficiary protection, and regulatory defensibility.

Exhibit 5 illustrates how SoE-led workflow execution redefines the claims process from a sequential model to an orchestrated workflow.

Exhibit 5: Before vs. after: SoE-led workflow execution (claims)

Source: Everest Group (2026)

Before SoE implementation	After SoE implementation
FNOL captured	FNOL triggers SoE
Claims wait in queue	Agent structures submissions
Evidence gathering	Coverage and severity routing begins early in the process
Vendors scheduled late	Fraud and anomaly signals run in parallel
Downstream fraud checks	Evidence requests and vendor scheduling orchestrated immediately after conditions are met
Settlement preparation basis adjuster's availability	Payments progress only through authority gates with SoD and audit logs
Payment action split across tools and handoffs	Customer communication stays synchronized to case state

Although SoE make the workflow orchestrated and autonomous through agents, the requirement of HITL remains crucial to achieve the desired results.

Agent vs. human: staged autonomy and mandatory HITL

SoE enable graduated autonomy rather than a binary agent versus human model, performing tasks such as:

- **Assistance:** documents/notes summarization, facts extraction, and customer updates drafting
- **Recommendations:** severity routing, evidence requests, next steps, and settlement ranges (with rationale)
- **Execution with approvals:** vendor scheduling, evidence requests orchestration, settlement packages preparation, and initiate payment steps under thresholds (approval required above thresholds and for sensitive decisions)
- **Execution with audit only:** low-complexity, rule-bound claims with complete evidence, backed by sampling and exception reporting

HITL is typically mandatory for denials/coverage disputes, high severity or injury claims, litigation-prone scenarios, suspicious claims requiring investigation judgment, and high-value payments.

SoE components required for this to work include:

- **Orchestration** for parallel workstreams (coverage, severity, fraud, vendors, and evidence)
- **Controls by design** for payment authority, duty segregation, audit trails, and evidence capture
- **Observability** for traceability, drift monitoring, exception patterns, and control adherence reporting

Value hypothesis: credible from day one, net of run costs

Claims value comes from compressing waiting time, reducing handling effort for routine coordination, and improving leakage control through earlier, consistent guardrails. A credible case includes:

- **Gross value:** cycle-time reduction, lower handling cost, improved customer communication cadence, and reduced leakage through disciplined controls
- **Cost to run:** HITL oversight and sampling, governance documentation and reviews, monitoring/evaluation, and integration maintenance

While early use cases demonstrate SoE's operational value, carriers must still determine how to scale these capabilities across the enterprise.

An SoE adoption roadmap

Once the need for SoE is realized, considerations are addressed and value chain elements selected, the next step is adoption. Transitioning to a SoE-centric model requires a structured, multi-phase approach that aligns technical maturity with business readiness.

Exhibit 6 presents a practical roadmap to transition from isolated pilots to enterprise-scale SoE.

Exhibit 6: SoE adoption roadmap

Source: Everest Group (2026)

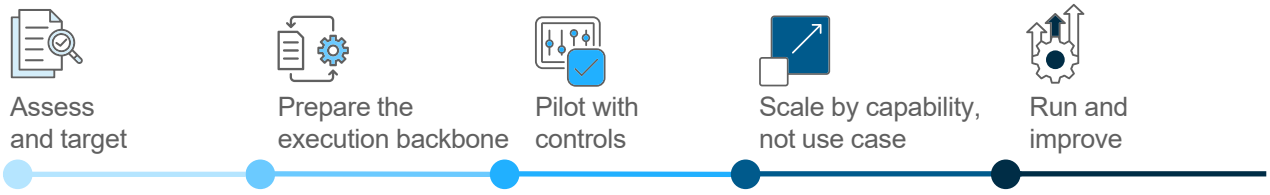

	Assess and target	Prepare the execution backbone	Pilot with controls	Scale by capability, not use case	Run and improve
Objective	Define the Minimum Viable Context (MVC) for the first use case: the smallest slice of data and policy needed for an agent to make a safe and defensible decision, where actions are frequent	Ensure AI agents have specific identities tied to access control, mirroring the authority level of human underwriters and adjusters	Constrain scope to a single domain and build evidence capture for every decision to ensure auditability from day one	Leverage reusable components for various business areas by standardizing patterns such as agent toolkits, governance templates, and integration adapters	Establish a continuous learning loop where agent performance is constantly evaluated against drift, policy adherence, and customer feedback through human override analytics
Key action	Identify manual tasks that are costing the most time and money, focusing on where handoffs are currently slowing the business down	Establish the context fabric, integrating real-time data from internal systems and third-party providers	Define the thresholds for human intervention based on confidence scores and risk levels	Transition from isolated projects to an enterprise-wide SoE framework	Develop regulatory packs to automate compliance reporting for state regulators

The next section highlights a case study that depicts SoE in action and the outcomes driven by its implementation and adoption.

Case study: SoE in action

Challenge

MSIG USA, a growing specialty insurer, wanted to modernize its claims workflows. As it scaled across complex risks, claims intake emerged as an opportunity to improve operational responsiveness. Intake relied heavily on manual interpretation of unstructured inputs across emails, documents, and internal systems, delaying the creation of decision-ready context. While core platforms supported transaction processing, MSIG USA sought to interpret data semantically and surface critical signals earlier in the workflow. First Notice of Loss (FNOL) timelines were longer than necessary, highlighting the need to move from data capture to execution readiness at intake.

Impact

The initiative significantly improved both speed and quality of intake as FNOL timelines reduced to under an hour, while data accuracy exceeded 90% across lines of business.

Early identification of critical claims enabled faster intervention. MSIG USA is now progressing toward same-day FNOL as a standard, enhancing customer responsiveness. More importantly, by establishing a structured, decision-ready context at intake, the insurer has created a scalable foundation for extending agentic AI across the claims life cycle and advancing toward SoE.

Solution

The company used NTT DATA's Agentic Insurance Workplace platform to introduce an agentic AI-led intake layer to transform how claims are understood and activated at entry. The solution focused on four key areas: structuring unstructured intake data, creating a consistent and decision-ready view of each claim, enabling earlier identification and prioritization based on key characteristics, and embedding governance through human oversight aligned with operational thresholds. NTT DATA delivered the transformation through an iterative, co-creation model, with rapid deployment and phased scaling across lines of business.

Constraints that will slow autonomy

While the adoption roadmap provides a structured path to enterprise-scale execution, carriers must acknowledge the structural constraints that influence how quickly and safely autonomy can be introduced. In Exhibit 7, we take a look at the typical constraints organizations will face in the path to execution and why they must address them.

Exhibit 7: Structural constraints shaping the design of SoE in insurance

Source: Everest Group (2026)

Constraint	What to track
Execution variability across insurance products and jurisdictions	Insurance workflows vary significantly by product type, jurisdiction, and operating model. Autonomous systems must adapt to these differences while maintaining consistent governance.
Decision-ready context is rarely available	AI agents require a structured, complete context to execute decisions reliably. However, insurance data is often fragmented across systems and unstructured documentation.
Legacy architecture constrains agentic execution	Most insurance core systems were designed for transaction processing rather than orchestrating AI-driven workflows. This limits what agents can execute directly.
Human oversight remains a permanent control layer	Insurance decisions often carry financial, legal, and regulatory implications. Human oversight remains essential for high-impact or sensitive decisions.
Governance and compliance must be scaled with autonomy	As agentic systems perform more decisions and actions, carriers must demonstrate auditability, policy compliance, and regulatory alignment at scale.

Once these constraints are understood, the next step is to translate their implications into SoE design. In Exhibit 8, we outline these implications for different insurance streams.

Exhibit 8: From constraint to design: SoE implications across P&C and L&A insurance workflows

Source: Everest Group (2026)

Constraint	P&C insurance implications	L&A insurance implications
Execution variability across insurance products and jurisdictions 	SoE must support configurable workflows and authority thresholds to accommodate varying claims handling rules and catastrophe scenarios Execution paths must integrate vendor ecosystems within governed orchestration flows	SoE must embed configurable suitability rules, disclosures, and regulatory checkpoints within underwriting and servicing workflows Execution flows must accommodate documentation-heavy processes such as medical evidence collection and beneficiary servicing
Decision-ready context is rarely available 	SoE must operationalize MVC gates during submission intake and claims processing to detect missing or inconsistent risk data Execution flows must trigger automated enrichment from external data sources	SoE must orchestrate structured evidence collection for underwriting context, including medical and financial documentation Workflow gates must ensure suitability, and policyholder information should be validated before autonomous decisions progress
Legacy architecture constrains agentic execution 	SoE must externalize orchestration outside core claims and policy systems to enable event-driven workflow progression Integration adapters must provide governed access to core system actions and vendor platforms	SoE must overlay orchestration across policy administration, underwriting, and servicing platforms to coordinate execution without heavy core modifications Integration layers must unify data flows across distribution, advisor, and servicing systems
Human oversight remains a permanent control layer 	SoE must embed graduated autonomy models with HITL checkpoints for high-value claims decisions and fraud triggers Authority thresholds and approval gates must govern settlement actions and coverage determinations	SoE must incorporate mandatory human checkpoints for underwriting exceptions, suitability reviews, and beneficiary claims Execution flows must capture evidence and rationale to support expert review and regulatory defensibility
Governance and compliance must be scaled with autonomy 	SoE must standardize control blueprints for claims payments, approvals, and exception handling Execution monitoring must capture decision logs, approvals, and audit trails across automated claims workflows	SoE must generate structured evidence for suitability decisions, beneficiary claims, and servicing actions Governance frameworks must enable automated compliance reporting and explainability for regulatory oversight

Addressing these constraints is necessary but not sufficient. Carriers must also ensure that execution improvements are measurable and sustained over time.

Exhibit 9 outlines a set of workflow-level metrics to track execution performance, control quality, and sustained value realization.

Exhibit 9: Key metrics to monitor at a workflow level

Source: Everest Group (2026)

Metrics	What to track
Execution outcomes	End-to-end cycle time (such as submission to quote and FNOL to settlement) Straight-through processing rate Rework rate (cases returning due to missing context or errors) Queue wait time versus processing time to expose orchestration tax
Decision and control quality	Exception rate and drivers Human override rate and rationale categories Control adherence (approval gates triggered as expected)
Risk and reliability	Drift indicators, including changes in agent behavior over time Escalation rates Error patterns by tool/integration where execution fails operationally
Cost of execution	Cost to run (HITL oversight load, monitoring, governance routines, integration maintenance, and third-party data costs) Cost to serve improvements Capacity released in expert roles and time shifted from coordination to judgment

Segment-specific indicators should also be monitored. In P&C insurance, value realization often appears through metrics such as leakage reduction, vendor cycle compression, and catastrophe response acceleration. In L&A insurance, improvements may manifest in underwriting cycle-time reduction, placement ratio improvement, servicing accuracy, and faster beneficiary claim settlement.

Ultimately, carriers that succeed with SoE will not simply deploy new technology, but will institutionalize new operating behaviors around execution.

Conclusion

The transition from a reactive carrier to an agentic carrier is not an overnight shift but a strategic migration that SoE enable. Successful organizations will be defined by their ability to bridge the gap between AI-driven insights and their operational realities.

Four critical approaches will differentiate successful SoE adopters from those stuck in pilots:

Build systems, not just agents

Treat agentic AI as operational infrastructure rather than a collection of isolated tools

Standardize capability

Build reusable patterns for ingestion, authority control, and legacy integration to enable scalable, consistent execution

Embrace the forward-deployed mindset

Break down technical silos by embedding cross-functional engineering teams directly into business units

Value transparency as much as efficiency

Bake governance and auditability into the core architecture, ensuring every autonomous action is defensible to regulators and customers alike

SoE are the industrial kitchen for modern carriers – a holistic system of ontologies, information flows, and technologies that ensure continuous interplay between automated software and human-led services. For carriers, the decision now centers on whether to continue layering point tools on brittle foundations or to build the execution architecture required to turn agentic intelligence into sustainable enterprise value.

Everest Group®
With you on the journey

Everest Group is a leading research firm helping business leaders make confident decisions. We guide clients through today's market challenges and strengthen their strategies by applying contextualized problem-solving to their unique situations. This drives maximized operational and financial performance and transformative experiences. Our deep expertise and tenacious research focused on technology, business processes, and engineering through the lenses of talent, sustainability, and sourcing delivers precise and action-oriented guidance. Find further details and in-depth content at www.everestgrp.com.

This study was funded, in part, by NTT DATA

For more information about Everest Group, please contact:

+1-214-451-3000

info@everestgrp.com

For more information about this topic please contact the author(s):

Aaditya Jain, Partner
aaditya.jain@everestgrp.com

Chinmay Pathak, Practice Director
chinmay.pathak@everestgrp.com

Rugved Sawant, Practice Director
rugved.sawant@everestgrp.com

Sohit Kumra, Senior Analyst
sohit.kumra@everestgrp.com

Notice and Disclaimers