Agentic AI in UAE motor claims

Preparing insurance carriers for compliant AI automations

Agentic AI in UAE motor claims

A whitepaper for UAE motor insurance leadership in response to the 4 May 2026 Dubai directive on agentic AI · Published May 2026 · Axxion Claims Settlement Services L.L.C.

1. Executive summary

On 4 May 2026, His Highness Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum directed Dubai's private sector to adopt agentic AI within 24 months. The directive followed a federal target announced 23 April 2026 to deliver 50% of UAE government services through autonomous AI agents by 2028. For UAE motor insurers, the political signal arrived ahead of any regulatory text. Boards now expect a response. Most operations cannot defend one yet.

The argument of this paper is direct: the 24 months the directive opens are operational rebuild time. An agent dropped on top of a claims operation that cannot reconstruct last week's decisions does not produce autonomy; it produces speed without traceability, which is a worse legal position than the one the carrier started with. The UAE Personal Data Protection Law (Federal Decree-Law 45/2021) still gives policyholders the right to object to fully automated decisions. The New CBUAE Law (Federal Decree-Law 6/2025) still keeps the licensed insurer accountable for every customer outcome, regardless of which system made the call. Compliance does not relax because the political ceiling has lifted.

The paper presents two arguments and three supporting frames:

- **Argument one (definitional).** What "self-executing" means in motor claims, stage by stage. The seven stages of a motor claim do not have a single agentic-eligibility answer. Some stages are honestly automatable today. Some are not, by law and by good operating sense. The conflation of all seven into a single "deploy AI" decision is where most procurement fails.
- **Argument two (diagnostic).** Six tests an insurer should run on its own operation before any agentic AI vendor enters the room. The tests examine whether the operation underneath has the data, the decision documentation, the licensing footprint, and the audit infrastructure to support an autonomous decision the regulator can later challenge.
- **Three supporting frames.** The seven-layer compliance stack that has to exist beneath any defensible deployment; the 24-month phased roadmap (data foundation, shadow-mode pilot, graduated autonomy); and the operational trap most insurers will walk into if they prioritize speed over record.

Axxion Claims Settlement Services L.L.C. operates the UAE's first dedicated motor claims operating system on these design principles. The operating shape was not built for the agentic AI conversation. It was built to be defensible under the New CBUAE Law and the PDPL. The shape required by the directive happens to be the same.

2. The mandate, in context

On Sunday 3 May 2026, with public announcement the following day, His Highness Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, Crown Prince of Dubai, directed the Emirate's private sector to adopt agentic AI within 24 months. The directive defines agentic AI as "self-executing and self-leading artificial intelligence" capable of analyzing, deciding, and acting with minimal human input [Government of Dubai Media Office, 4 May 2026]. The directive operationalizes through Dubai Chambers, with three mechanisms: training tracks for all business councils, incubators for agentic AI companies, and dedicated investment funds.

The Dubai-level directive sits on top of a federal target announced 23 April 2026: 50% of UAE government services delivered through autonomous AI agents by 2028, with implementation under HH Sheikh Mansour bin Zayed Al Nahyan and the policy task force led by HE Mohammed Al Gergawi [Khaleej Times; Gulf News].

What changed between 22 April and 5 May 2026 is the political ceiling. What did not change is the regulatory floor. The Dubai Universal Blueprint for AI remains the existing policy frame. The New CBUAE Law (Federal Decree-Law 6/2025) entered force in 2025 with a one-year reconciliation window, consolidating supervision of insurance and financial activity under the Central Bank. The UAE Personal Data Protection Law (Federal Decree-Law 45/2021) continues to govern the rights of data subjects, including the right to object to decisions made entirely by automated processing where those decisions produce legal or similarly significant effects. A motor claim outcome of accept, reject, total loss, or supplement produces such an effect.

The mandate operates as a political directive without statutory force. There is no published compliance test, no certification, no binding penalty regime tied to it. The legal framework that any agentic AI deployment must sit inside has not relaxed in response. The question for the UAE motor insurer is therefore what the operation underneath can defend when an autonomous decision goes wrong.

The 24-month deadline may extend. Political directives have moved before, and the private sector scope is softer than the federal target. The argument of this paper does not depend on the deadline holding to the day. The political pressure persists even if the timeline drifts; the legal framework continues to apply regardless of when the directive is judged complete; and the operational rebuild proposed below produces a defensible motor claims operation whether the deadline lands at 18 months, 24 months, or 30. Insurers waiting for timeline certainty before starting are choosing to lose either way: deadline holds and they are late, deadline slips and they have wasted the runway.

3. What "self-executing" means in UAE motor claims

Agentic AI reads in a press release as a single capability. A motor claims operation is a sequence of distinct decisions, each with its own evidence base, regulatory source, and customer-impact profile. The first analytical move for any insurer responding to the mandate is to refuse the conflation.

The Axxion Claims OS organizes the motor claims lifecycle into seven sequential stages, each with its own compliance gates and decision logic. Comparing each stage against the agentic-eligibility criterion (whether a decision can be made autonomously without breaching law or governance) produces the only true answer to "where does agentic AI fit?"

Stage	Agentic eligibility	Operational and regulatory boundary
1. FNOL and intake	Predominantly eligible today	Notification capture, claim type classification, policy validation, file opening, structured documentation request, and severity triage are activities vendor platforms perform at scale internationally. The constraint is data quality at intake.
2. Triage and assessment	Partially eligible	Damage classification from photographs, repair-versus-total-loss recommendation, and exception flagging are automatable. The boundary is the loss adjusting line. Surveying and Loss Adjusting is a CBUAE-regulated insurance-related profession; an autonomous system that determines damage assessments and approves repair scopes is conducting loss adjusting and must sit inside an entity holding the relevant CBUAE registration.
3. Workshop allocation	Eligible with rule-based governance	Routing a claim to the correct workshop based on vehicle type, geography, capability assessment, and current capacity is structured decisioning. Operational risk sits upstream, in whether the workshop network is governed at all.
4. Repair management	Mixed	Estimate review against benchmarks, supplement triage, milestone tracking, and parts authorization can carry significant automation. Workshop quality intervention, structural repair sign-off, and OEM-mandated calibrations require licensed human oversight. The default is graduated autonomy: agents propose, humans approve at defined thresholds, with reasoning recorded for both.
5. Quality verification	Limited	Computer vision can pre-screen for paint match, panel gap, and obvious defects. Final quality sign-off attaches to a named accountable human, both because the customer expects it and because the audit standard for any defect dispute requires an individual identity in the record.
6. Vehicle handback and customer communication	Highly eligible for status updates; constrained for adverse decisions	Status updates and structured communications are well suited to automation. Adverse-decision communication is constrained: under the PDPL, a policyholder receiving a fully automated decision with significant effect has explicit rights of objection and explanation. The communication channel must preserve those rights at the point of decision delivery.
7. Settlement and recovery	Partially eligible	Reconciliation against the original estimate, payment scheduling, and recovery file initiation are automatable. The AML/CFT gates (sanctions screening, beneficiary verification, suspicious-activity flagging) can be performed by automated systems, but the licensed entity remains accountable for failures regardless of which component made the call [Federal Decree-Law 20/2018 on AML/CFT; CBUAE AML/CFT framework].

Most stages permit material automation. None permits unsupervised end-to-end autonomy without breaching at least one regulatory or contractual constraint today. The useful question for an insurer is which decisions in the claim can be agentic, under what governance, and recorded how. That question has stage-level answers, with the platform choice flowing from the answer rather than determining it.

> An insurer that opens with the vendor selection question is two steps too late. The first question is which stages of the claim the regulatory perimeter permits autonomy in, and what governance has to wrap the rest. The carrier's operating model holds that answer; vendors cannot generate it from outside. Carriers that resolve the stage-by-stage question before procurement save twelve to eighteen months of integration rework.

4. Six tests for an agentic-AI-ready motor claims operation

Before any vendor enters the building, an insurer's leadership team should apply six tests to its own operation. Each test has a binary answer. Each examines an operational reality at the carrier today. The operation that passes the six tests can support an agentic AI deployment a regulator will accept. The operation that fails three or more should run an operational rebuild before a procurement decision.

The tests are derived from the legal framework, the regulatory direction of travel, and the operating signature of working agentic systems internationally. They reflect a particular operational philosophy: compliance-first, evidence-led, audit-ready. An insurer working from a different philosophy may legitimately reach different conclusions on what "agentic-AI-ready" requires. The diagnostic above is most useful for insurers whose answer to "what does the regulator need to see" is "everything, on demand."

An insurer that completes the six knows where to start. An insurer that defers them ends up choosing a platform on the strength of its demo, rather than on the strength of its fit with the carrier's actual data foundation.

Test 1: Reconstructable decisions

Pick six claims at random from last week. Try to reconstruct, for each, the full decision history: the inputs received, the assessments made, the reasoning for each authorization, the reviewer where required, the final cost reconciliation. Allow one working day to assemble each.

The pass standard is six of six, with the trail produced from the system of record. Individual emails, WhatsApp histories, and memory do not satisfy the standard. Most UAE motor insurers cannot pass this test on three of six today, and that estimate is generous. The decision history sits in someone's email, the photo evidence sits in a shared folder, and the reasoning lives in the head of the senior claims handler who reviewed the file.

A regulator running the same exercise on agentic-AI-decided claims, with eighteen months' notice and several hundred files instead of six, expects the same standard. The system that produces the decisions also produces the trail. Either the operation is designed that way today, or the agentic AI deployment is built on a foundation that cannot defend itself in a CBUAE inspection.

If the operation cannot pass this test, the data foundation is the first project and the AI is the second, in that order.

Test 2: Structured reason codes for every outcome

Examine claim outcomes in the system. The relevant question is whether decisions sit in structured reason codes (a finite, governed taxonomy a system can read) or in free-text notes only a human can interpret.

Reason codes are the language an agent reads; free text registers as noise. An automated system cannot reliably learn from "approved per discussion with WS" or "client called, OK to release"; it cannot benchmark against historical patterns when the patterns sit inside unparseable strings. The downstream consequence is that pre-deployment testing of any agentic AI system becomes an exercise in retrofitting structure onto unstructured history, rather than supervised learning against a clean ground truth.

The pass standard is that every claim outcome (accept, reject, total loss, supplement approval, supplement rejection, recovery filed, recovery declined) attaches to a structured code with mandatory metadata: reason category, sub-reason, financial impact, timestamp, identity of decider. The free-text field exists for context. The decision sits in the structured code.

Test 3: Data residency and PDPL transfer mechanism

For each component of the proposed agentic AI architecture, identify where the personal data physically resides during processing, which legal entity is the data processor at each hop, and under what PDPL cross-border transfer mechanism the data moves between components.

Most off-the-shelf agentic AI platforms run on hyperscaler infrastructure outside the UAE. Personal data flowing to those platforms is a cross-border transfer. The PDPL imposes restrictions: transfer is permitted to jurisdictions with adequate protection, or with appropriate safeguards documented in the data processing arrangement [UAE PDPL, FDL 45/2021; UAE Data Office cross-border transfer guidance]. Claims data (vehicle identifiers, policyholder identity, accident circumstances, medical detail where injury is involved, financial detail) is personal data. An imported model usually means an imported data flow. Properly safeguarded hyperscaler deployments remain PDPL-permissible. The compliance failure point is undocumented transfer. Without a documented mechanism, the deployment is non-compliant from day one.

The pass standard is a data flow diagram showing every hop, an identified PDPL transfer basis for each hop, and a contractual chain that holds every processor to the same standard.

Test 4: Licensing scope and the loss adjusting line

For each decision the agentic system is proposed to make autonomously, identify which CBUAE-regulated activity it implicates and which licensed entity in the corporate structure is the accountable performer.

This is the test most platform demos cannot answer. Loss adjusting, defined as the assessment of damage and the determination of the cost basis, is a regulated insurance-related profession under the New CBUAE Law. An autonomous system performing loss adjusting decisions, deployed in-house by an insurer without the regulated capability sitting inside a properly licensed entity, is operating outside the licensing perimeter regardless of how good the underlying technology is.

The pass standard is a clear allocation: each agentic decision sits inside the activity scope of an entity that holds the relevant CBUAE registration, with the human accountable individual named for the decisions that require one.

Test 5: Complaint flow under FDL 6/2025

Trace what happens when a policyholder contests an automated decision. The customer files a complaint at the carrier. Within the conduct-of-business framework operating under the New CBUAE Law, the complaint must be acknowledged, investigated, and resolved within defined timelines, with a documented audit trail [New CBUAE Law, FDL 6/2025, customer protection and complaints provisions]. If unresolved, the complaint can escalate to Sanadak, the UAE's independent insurance and financial ombudsman, and ultimately to the Insurance Disputes Settlement and Resolution Committee.

For each tier, the carrier must produce: the original automated decision, the inputs that produced it, the reasoning the system applied, the model version in operation at the time, the human checkpoint where required, and the final resolution. If any element is missing, the carrier loses the dispute by default. The two elements that fail most often are the model version and the inputs at the time of decision. A favorable underlying claim outcome does not save a process that cannot evidence itself.

The pass standard is a runbook that walks an inspector from complaint receipt to dispute resolution with every artifact accessible within working hours. Most carriers cannot produce this runbook for human-made decisions. The standard does not relax for AI-made ones.

Test 6: Vendor liability and indemnity

Read the proposed vendor agreement. Locate the indemnity clauses, the warranty on accuracy, the audit rights, and the carve-outs for AI-specific failure modes (hallucination, drift, training-data bias, model versioning).

When the AI system makes a wrong decision, the licensed insurer is liable to the policyholder. Vendor agreements must allocate that risk back to the vendor where the failure mode is the vendor's responsibility. Most off-the-shelf vendor agreements do not. Standard SaaS terms cap liability at fees paid in the prior year, a number that will not survive a single wrongful claim denial at scale, let alone a class of them.

The pass standard is a vendor agreement with: SLAs on decision accuracy with measurable thresholds, indemnity sized to the regulatory exposure rather than the contract value, audit rights with reasonable notice, model versioning commitments with change-control approvals, and explicit allocation of liability for AI-specific failure modes.

Reading the test results

Score	Reading	Recommended next step
6 of 6	The operation can support agentic AI procurement.	Begin shadow-mode pilot. The mandate timeline is workable.
5 of 6	One specific gap.	Identify and fix the gap. Defer the AI conversation by 60 to 90 days. Procurement starts with the gap closed.
3 or 4 of 6	Multiple foundation gaps.	The operation needs an operational rebuild before any platform conversation. An agent layered onto this foundation amplifies weaknesses faster than it delivers benefit.
2 or fewer	The agentic AI conversation is two phases ahead of where the operation sits today.	Build the foundation first. The 24-month window is achievable but only with a sequenced program; vendor procurement in months one to six is the wrong move.

> The mandate creates a board-paper deadline. The six tests create the operational deadline. Insurers that schedule a vendor demo before scoring their own operation against the six are buying a platform that will land on a foundation it cannot defend. Running the six tests internally takes two to three weeks. The cost of skipping them is an eighteen-month integration that ends in a regulatory finding.

5. The compliance stack underneath every defensible deployment

The six tests are diagnostic. Behind them sits a structural answer: the seven-layer compliance stack that any agentic AI deployment in motor claims has to be built on. Each layer corresponds to a specific UAE legal source. Each is independently necessary. None is sufficient on its own.

- **Data architecture.** Structured input data, canonical identifiers, single-source-of-truth discipline, complete event capture. The PDPL basis is lawful processing and transparency. The operational test is reason codes (Test 2).
- **Decision documentation.** Every autonomous decision recorded with inputs, model version, decision logic, confidence score, reason code, and final action. The CBUAE source is conduct-of-business recordkeeping. The operational test is reconstructability (Test 1).
- **Model operations.** Versioned models, controlled training data, drift monitoring, change management. The basis is the supervisory expectation that applies to any algorithmic decisioning regulated under the Central Bank's prudential framework.
- **Human governance.** Defined human checkpoints, threshold-based escalation, named accountable individuals for decisions that require them. The source is the governance and management responsibility framework under the New CBUAE Law.
- **Complaint handling.** A complaint pathway that preserves access to the original decision artifacts, supports objection rights under PDPL, and meets the conduct-of-business timelines under FDL 6/2025 (Test 5).
- **Audit infrastructure.** A data and decision trail an inspector can query within working hours, end-to-end across the seven-stage pipeline. The basis is the supervisory expectation that the entity can evidence what it has done. Attestation alone falls short of that standard.
- **Regulatory licensing scope.** Each agentic decision sits inside the activity scope of an entity that holds the relevant CBUAE registration, with no autonomous loss adjusting outside a licensed perimeter (Test 4).

The seven layers are operational infrastructure that the carrier owns. A vendor demo that promises layers four through six "out of the box" is selling capabilities that have to be customized into the carrier's actual operating model. Most of the work belongs to the carrier. Procurement decisions that assume otherwise produce the integration overruns that have already become public in the UK and Australian motor markets where similar deployments ran ahead of their operating foundation.

6. The 24-month roadmap: three phases

The 24-month window the directive opens fits a phased operational rebuild. The phases are sequential by design. A phase-two pilot on a phase-one foundation is testable. A phase-two pilot without that foundation is theater.

Phase	Months	Focus and deliverable
Phase 1 Foundation	0–6	Restructure claim outcomes into a governed taxonomy. Stand up canonical identifiers. Document data flows for PDPL compliance. Compare the seven-stage pipeline against current operations and identify gate gaps. Fix the audit trail breaks. The deliverable at month six is the ability to pass Tests 1 and 2 with evidence.
Phase 2 Shadow mode	6–18	Deploy the proposed agentic AI in observe-only mode against live claim flow. The system produces its decisions; humans make the actual decisions; both are recorded. Compare. Calibrate. Identify failure modes the demo did not surface. The deliverable at month eighteen is a comparison dataset large enough, working estimate of 5,000 to 10,000 claims, to establish accuracy and confidence-score reliability against operational ground truth.
Phase 3 Graduated autonomy	18–24	Where the shadow-mode pilot demonstrates accuracy at or above the human baseline, with confidence scoring that correctly identifies its own errors, transition specific decision types to agentic authority. Maintain human checkpoints for decisions that touch the loss adjusting line, the complaint flow, and the AML gates. Run the operation in graduated mode for the remainder of the window, reviewing escalation triggers monthly.

The roadmap assumes the corporate structure question is resolved separately. An insurer carrying out autonomous loss adjusting decisions in-house, without those decisions sitting inside a properly licensed entity, has a structural problem no roadmap phase fixes. The licensing footprint is a precondition that has to be resolved before any phase begins.

> An insurer that signs an agentic AI procurement in May 2026 with a 12-month implementation target is committing to deploy on a foundation it has not yet built. The right schedule is build foundation, then sign procurement, then deploy. Reversing that order has produced more regulatory findings in adjacent industries than any technology choice has.

7. The trap to avoid: speed without record

The mandate creates pressure. Pressure produces shortcuts. The shortcut most likely to be taken in the UAE motor insurance market over the next 24 months is to layer an agentic AI on top of the existing operation without rebuilding the data foundation underneath it.

The argument for the shortcut is intuitive. The directive is real, the timeline is short, the procurement budget is approved, and the vendor demo is impressive. The reasoning then runs that the foundation problems can be fixed in parallel: get the AI deployed, capture the operational benefits early, build the audit trail alongside.

The strongest version of the parallel-rebuild case is real and worth engaging. A vendor with a strong integration architecture, a carrier with internal data engineering capacity, and a program manager who treats the foundation work as an equal track rather than a subordinate one can sometimes deliver both in step. The test of legitimacy is whether the foundation rebuild has its own funded workstream, its own milestones, and its own deliverables that ship on time, or whether it becomes the work that gets deferred when the AI deployment hits its first integration issue. In the carriers where parallel rebuild has worked internationally, the foundation program had a director-level owner with veto power over the AI deployment timeline. In the carriers where it failed, the foundation work was assigned to whoever had spare capacity. The structural test is governance.

For carriers that cannot meet that governance standard, the reasoning fails on three points. First, an agent making decisions on undocumented inputs cannot be debugged after the fact. When a wrongful denial reaches Sanadak, the carrier cannot reconstruct what the system was looking at when it decided; the dispute defaults against the carrier. Second, retrofitting an audit trail to decisions already made is not feasible at scale. The decisions were made on data that was not captured; the trail is not buildable after the event. Third, the insurer becomes operationally dependent on a vendor whose platform cannot be replaced without losing the decision history embedded in it.

The shortcut produces speed in months one through six. It produces immobility from month seven onward. The carrier that took the shortcut cannot move to a different platform without losing institutional decision memory, cannot defend its existing decisions in regulatory dispute, and cannot upgrade the foundation without unwinding the platform sitting on top of it.

The carriers that resist the shortcut and rebuild the foundation first will be twelve to eighteen months slower to a deployment headline. They will have a deployment that survives the first regulatory inspection. The first inspection is the only test that matters.

8. Where Axxion fits

Axxion Claims Settlement Services L.L.C. operates the UAE's first dedicated motor claims operating system. The Claims OS architecture (six layers, seven stages, compliance-by-design) was not built as a response to the agentic AI directive. It was built to be defensible under the New CBUAE Law and the UAE PDPL. The operating shape required by the directive happens to be the same.

Three observations relevant to insurers running through the six tests:

- The Claims OS records every authorization decision with the input data, the logic applied, the reviewer where required, the reason code, and the final action. Tests 1 and 2 sit as operational defaults inside the platform today.
- Workshop allocation is rule-driven against a capability-assessed network. Settlement reconciliation runs against the original estimate, with the final invoice reviewed against that benchmark. Cross-border data flows operate under documented PDPL transfer mechanisms.
- The licensing footprint separates regulated activities (claims management performed for licensed insurers) from non-regulated activities (workshop network management, repair coordination, data services), keeping each function inside a properly authorized entity.

Insurers that score five or six on the six tests are positioned to procure agentic AI directly. Insurers that score three or four can either rebuild the foundation in-house (typically a 12 to 24 month program depending on the gap), or partner with an operating model already running to the standard. The motor line is one part of the enterprise mandate response. The rest of the enterprise program sits alongside it.

The most useful conversation with Axxion is a structured walkthrough of the six-test diagnostic against the carrier's current operation, compared with what an Axxion-supported motor claims engagement would change. The conversation runs at the operational level. The output is a written assessment of where the operation passes, where it fails, and what a realistic path to a defensible deployment looks like.

Reader next step. A 30-minute structured conversation walks the six-test diagnostic against the carrier's current operation. Direct contact: hi@axxion.co.

9. Sources

Primary mandate sources.

- Government of Dubai Media Office, 4 May 2026: mediaoffice.ae/en/news/2026/may/04-05/hamdan-bin-mohammed-launches-dubai-private-sector-shift-to-agentic-ai-within-two-years
- Khaleej Times: federal directive on UAE government services and AI agents, 23 April 2026
- Gulf News: federal directive coverage, 23 April 2026
- Dubai Universal Blueprint for AI (existing policy frame)

Regulatory framework.

- UAE Personal Data Protection Law: Federal Decree-Law No. 45 of 2021. Specific article references applicable to automated decisioning rights are described at the level of effect and right; specific article numbers should be confirmed against the official gazette before legal reliance.
- New CBUAE Law: Federal Decree-Law No. 6 of 2025 Regarding the Central Bank, Regulation of Financial Institutions and Activities, and Insurance Business. Article 185 repealed Federal Decree-Law No. 48 of 2023; Article 184 sets the one-year reconciliation window.
- UAE Federal Decree-Law No. 20 of 2018 on Anti-Money Laundering and Combating the Financing of Terrorism and Financing of Illegal Organisations (as amended).
- CBUAE Surveyor and Loss Adjuster regulatory perimeter.
- CBUAE customer protection and complaints framework; Sanadak ombudsman framework.

Industry and competitive sources.

- BCG: Agentic AI: Powering Core Insurance AI Modernization (2026).
- Norton Rose Fulbright: AI Innovation and Adoption in Insurance in the Middle East.
- InsurtechGulf: Agentic AI in Insurance Underwriting, 25 April 2026.
- Deloitte: State of Generative AI in the Enterprise.

Internal Axxion sources.

- Axxion Claims OS Blueprint, v0.5 (March 2026): six-layer architecture, seven-stage pipeline, compliance-by-design specification.
- Axxion strategy paper: Response to the Dubai agentic AI mandate, v0.2 (8 May 2026).
- Axxion LinkedIn series: Dubai agentic AI mandate, v0.5 (8 May 2026).

About Axxion

Axxion Claims Settlement Services L.L.C. is a Dubai-based motor claims management company and the UAE's first dedicated motor third-party administrator (TPA). Axxion manages the full motor claims lifecycle on behalf of insurance partners, from first notification of loss through damage assessment, repair coordination, quality control, and settlement. The operation pairs more than four decades of hands-on repair and motor claims expertise with AI-enabled processes to deliver lower repair costs, shorter cycle times, and auditable compliance on every claim.

Axxion's claims platform generates a documented cost trail on each claim, produces burning cost analytics for insurer partners.

The company is led by Managing Director and Co-founder Frederik Bisbjerg, an internationally recognized insurance executive whose career includes C-level leadership at Qatar Insurance Group, AXA Global Healthcare, Al Wathba Insurance, and Daman National Health Insurance. Bisbjerg is a published author on insurance transformation and a founding faculty member of the world's first mini-MBA in Digital Insurance.

His work as Head of MENA at The Digital Insurer and his contributions to AI strategy across the GCC have made him one of the region's leading voices on the application of artificial intelligence in insurance operations.

Axxion operates within World Automotive Group, a MENA-based automotive and insurance services group. World Automotive Group is owned by Skelmore Holding, a global consortium founded in Toronto in 1994, with $650 million in revenue and 4,000 employees across the GCC and North America.

Informational disclaimer. *This whitepaper is published for informational and discussion purposes only. It does not constitute legal advice, regulatory opinion, or investment recommendation. References to UAE Federal Decree-Laws, CBUAE provisions, and other regulatory instruments reflect Axxion's operational understanding as of 8 May 2026 and are not a substitute for qualified legal counsel. Insurers and other readers should consult their own legal and compliance advisors before making decisions on agentic AI deployment, vendor selection, corporate structure, or regulatory compliance strategy. The Central Bank of the UAE has not endorsed this paper and bears no responsibility for its content. Axxion Claims Settlement Services L.L.C. is a motor claims management company licensed in the UAE; the views expressed reflect operational experience in the motor claims sector and should not be read as advice on activities outside that scope.*