T H E M E A S U R E D V I E W

AIQA Global's series on the ideas, standards, and market forces shaping enterprise AI governance. In a field defined by speed and speculation, these articles offer something different: informed perspective grounded in the discipline of measurement. Because the organizations that will lead in AI are the ones that can prove it.

Key Person AI

What Happens to Your AI When Your AI Leader Walks Out the Door

AIQA Global, LLC
February 2026

The Person Who Knows

Every enterprise that uses AI has one. Sometimes it is the chief data scientist who built the first models and chose the architecture. Sometimes it is the machine learning engineer who understands the training pipeline, knows which data sources are reliable, and can explain why the model behaves the way it does. Sometimes it is not a technologist at all but a business leader who championed the AI initiative, secured the budget, negotiated the vendor contracts, and serves as the institutional bridge between the technical team and the board.

Whatever the title, the pattern is the same: a disproportionate share of the organization's AI knowledge, context, and decision history lives in the head of one person—or a very small number of people. In software engineering, this concentration is called the "bus factor": the number of people who would need to leave before a project stalls.[1] In many enterprise AI programs, the bus factor is likely close to one.

This is the key person risk of AI—and it carries dimensions that go beyond traditional key person risk. When a senior executive departs, the organization loses relationships, judgment, and institutional memory. When the person who understands your AI systems departs, the organization may lose the ability to explain, maintain, modify, or even operate the systems that increasingly drive its business decisions. The knowledge that walks out the door is not merely strategic. It is operational, technical, and—if ungoverned—irreplaceable.

Why AI Makes Key Person Risk Worse

Key person risk is not new. Organizations have always depended on individuals with specialized knowledge. But AI introduces characteristics that amplify this vulnerability in ways that traditional key person risk frameworks do not capture.

AI systems are opaque by nature. A traditional business process can generally be reverse-engineered from its documentation, its inputs, and its outputs. A machine learning model cannot. The relationship between a model's training data, its architecture, its hyperparameters, and its outputs is mathematically complex and often non-intuitive. The person who trained the model may understand why it performs well in some contexts and poorly in others. Without that person, the organization is left with a system it uses but does not fully understand—what practitioner Pavel Shpin characterizes as "the Black Box Whisperer" problem.[2]

AI knowledge is experiential, not just procedural. The most critical AI knowledge is often not the kind that can be written in a manual. It includes the data scientist's judgment about which features matter, the engineer's intuition about why a model drifts in certain conditions, and the architect's understanding of undocumented trade-offs made during development. This experiential knowledge—sometimes called tribal knowledge—accumulates over months or years of working with a specific system. It is precisely the kind of knowledge that is hardest to transfer and most damaging to lose.

AI talent is scarce and mobile. The market for AI and machine learning professionals remains exceptionally competitive. Deloitte's first quarter 2024 CFO Signals report found that sixty percent of surveyed chief financial officers said bringing in talent with generative AI skills over the next two years was either extremely or very important.[3] Yet the supply of experienced AI governance and engineering talent is constrained. When a key AI professional leaves, the replacement timeline is measured in months, not weeks—and during that interval, the organization's AI systems must still be maintained, monitored, and explained.

The cost of replacement is amplified. Gallup estimates that the cost of replacing a professional in a technical role is approximately eighty percent of that person's annual salary, with the cost for leaders and managers reaching two hundred percent.[4] For AI roles, the true cost may be higher, given the specialized knowledge transfer burden: the replacement must not only learn the job but reconstruct the contextual understanding of systems that may be poorly documented. The ramp-up period for an AI specialist inheriting undocumented models, pipelines, and decision logic can be measured in quarters, not weeks.

The Documentation Deficit

The root cause of AI key person risk is not talent scarcity. It is documentation failure. The available evidence suggests that enterprise AI systems are broadly underdocumented—not because organizations lack the resources, but because the culture of AI development has historically treated documentation as an afterthought.

Consider what comprehensive AI documentation would include: a description of the system's intended purpose and design specifications; the data sources used for training, validation, and testing, including their provenance and known limitations; the model architecture, algorithms, and hyperparameters selected, with rationale for design decisions; the testing and validation methodology, including performance metrics and their interpretation; a record of changes made to the system throughout its lifecycle; the risk management measures applied; and a plan for ongoing monitoring. This is not an aspirational wish list. It is, almost verbatim, the content required by Annex IV of the EU AI Act for every high-risk AI system.[5]

The gap between what regulation requires and what most organizations actually document is vast. IBM's 2025 Cost of a Data Breach Report found that sixty-three percent of organizations reported having no AI governance policies in place.[6] Organizations that lack formal governance policies are less likely to have the granular system-level documentation that would allow someone new to understand, maintain, and explain an existing AI system. The institutional knowledge is in the people, not in the records.

This creates a fragility that boards and investors are only beginning to appreciate. When the person who built and maintains your AI systems leaves, the question is not just "can we hire a replacement?" It is "can anyone explain what our AI does, why it does it, how it was built, what data it uses, what risks it carries, and how to fix it if it breaks?" If the answer depends entirely on a departing employee's memory, the organization has a governance problem that no amount of recruiting can solve quickly.

Regulation Demands What Key Persons Currently Carry

The regulatory landscape is making AI documentation an obligation, not a preference. The EU AI Act's technical documentation requirements under Article 11 and Annex IV are explicit: documentation must be prepared before a high-risk AI system is placed on the market, must be kept up to date throughout the system's lifecycle, and must be retained for at least ten years.[7] The documentation must include design specifications, training data descriptions, risk management systems, testing procedures, lifecycle change logs, and post-market monitoring plans.

These requirements are, in effect, a regulatory mandate for exactly the kind of institutional knowledge transfer that key person risk threatens to prevent. An organization that satisfies Annex IV has, by definition, externalized the critical AI knowledge that would otherwise reside solely in its people. An organization that does not has created a compliance gap that is also a business continuity gap.

The NIST AI Risk Management Framework reinforces this point through its Govern and Map functions, which call for documented policies, roles, accountability structures, and risk contextualization across the AI lifecycle.[8] ISO/IEC 42001, the international standard for AI management systems, addresses documented procedures for AI system lifecycle management, including policies, processes, and controls.[9] The convergence of these frameworks around documentation is not coincidental. It reflects a consensus that AI governance cannot depend on individual knowledge. It must be embedded in organizational systems and records.

Insurers Will Ask the Question

As AI-specific insurance products develop, underwriters are likely to evaluate key person risk as a component of AI governance quality. While this practice is not yet widespread for AI specifically, the logic follows established underwriting principles: an organization whose critical systems depend on undocumented knowledge held by a small number of individuals presents a higher risk profile than one whose systems are fully documented, transferable, and independently auditable.

The analogy to existing practice is direct. Professional liability underwriters already assess whether a firm's client relationships are institutionalized or concentrated in individual partners. Cyber insurers already evaluate whether an organization's security posture depends on a single administrator or is distributed across a team with documented procedures. It is reasonable to expect that AI governance quality assessment will follow the same pattern: documentation breadth, knowledge distribution, succession planning, and the ability to maintain and explain AI systems in the absence of any single individual are likely to become factors in governance scoring and, ultimately, in coverage terms.

The question an insurer will eventually ask is not "do you have AI systems?" but "if your lead AI engineer resigned tomorrow, could you still explain to a regulator how your AI makes decisions, demonstrate that it was tested for bias, and produce the documentation trail showing how it has been maintained since deployment?" For most enterprises today, the honest answer would be no. That answer carries underwriting consequences that are only beginning to materialize.

Building AI Resilience: From Persons to Processes

Addressing AI key person risk requires a deliberate shift from person-dependent to process-dependent AI operations. This is not a technology problem. It is a governance and management problem that requires investment in four areas.

Document everything, continuously. AI documentation should not be a one-time compliance exercise. It should be a continuous practice integrated into the development and deployment workflow. Every model decision, architecture choice, data source selection, and parameter adjustment should be recorded as it happens—not reconstructed after the fact. The EU AI Act's Annex IV provides a practical template for the minimum documentation required. Organizations that adopt this standard voluntarily, before regulatory enforcement compels it, will build institutional knowledge resilience as a byproduct of compliance.

Cross-train deliberately. No critical AI system should be understood by only one person. Cross-training is not a luxury; it is a risk management imperative. This means ensuring that multiple team members understand each system's architecture, data dependencies, performance characteristics, and failure modes. Pair programming, code reviews, and internal knowledge-sharing sessions are operational investments that raise the bus factor from one to something more resilient.

Separate the builder from the maintainer. Organizations that allow the same person who built an AI system to serve as its sole operator, monitor, and explainer have created a single point of failure by design. Governance frameworks should require that AI systems be transferable—that the documentation and operational procedures are sufficient for someone other than the original developer to maintain and monitor the system. If a system cannot be operated by someone who did not build it, it is not governed; it is merely staffed.

Assess and score documentation quality. Documentation quality should be measured, not assumed. Standardized governance assessments that evaluate documentation breadth, accuracy, and accessibility provide organizations with a quantifiable indicator of their resilience to key person departures. This is one of the governance dimensions that AI governance ratings are designed to capture: the extent to which an organization's AI systems are documented, transferable, and independently auditable—regardless of which individuals are employed at any given time.

Your AI Is Only as Durable as Its Documentation

The value of an enterprise's AI systems is not measured only by what those systems produce today. It is measured by whether the organization can maintain, explain, improve, and

defend those systems tomorrow—regardless of who is still on the payroll. An AI system that produces excellent results but cannot be understood, maintained, or explained by anyone other than its creator is not an asset. It is a liability with a timer on it.

Key person risk in AI is not an edge case. It is a common condition across enterprises, because many organizations have not yet built the documentation infrastructure, cross-training practices, and governance frameworks that would make their AI systems resilient to personnel changes. The organizations that address this now—by documenting their systems, distributing their knowledge, and measuring their governance quality—will be the ones that can scale AI with confidence. The rest will discover, sooner or later, that the most important thing about their AI walked out the door.

The irony is that addressing key person AI risk does not require heroic measures. It requires the same governance discipline that organizations already apply to financial controls, cybersecurity, and regulatory compliance: document the process, distribute the knowledge, assign accountability, and measure the result. The tools exist. The regulatory frameworks demand it. The insurance market is moving in the same direction. The only question is whether your organization will build this resilience proactively—or discover the gap at the worst possible moment.

Notes

1. César Soto-Valero, "Bus Factor: A Human-Centered Risk Metric in the Software Supply Chain," February 2022. Soto-Valero defines the bus factor as a measure of a project's resilience to the most extreme personnel losses, related to socio-technical debt caused by knowledge concentration. For a practitioner-oriented discussion of its application in the AI era, see also Pavel Shpin, "Your 10x Engineer Might Be a 'Bus Factor' of 1: De-risking Key Person Dependency," MindCTO, August 2025.

2. Shpin, "Your 10x Engineer Might Be a 'Bus Factor' of 1," MindCTO, August 2025. Shpin characterizes the "Black Box Whisperer" as the data scientist or engineer who has developed intuitive understanding of a non-deterministic model's behavior—knowledge that is experiential, not procedural, and effectively irreplaceable when the individual leaves. Note: This is Shpin's framing, not a standard industry term of art.

3. Deloitte, "CFO Signals: What North America's Top Finance Executives Are Thinking and Doing," Q1 2024. Sixty percent of surveyed chief financial officers stated that bringing in talent with generative AI skills over the next two years was extremely or very important. See https://www.deloitte.com/us/en/programs/chief-financial-officer/articles/cfo-signals-1q-2024.html. Also cited in Harvard Law School Forum on Corporate Governance, "Largest Companies View AI as a Risk Multiplier," November 20, 2024.

4. Gallup, "42% of Employee Turnover Is Preventable but Often Ignored," Gallup Workplace, July 10, 2024 (updated February 16, 2026). Based on a nationally representative survey of 717 U.S. employees who voluntarily left an employer within the past year. Gallup estimates replacement costs of approximately 200% of annual salary for leaders and managers, 80% for technical professionals, and 40% for frontline workers. See https://www.gallup.com/workplace/646538/employee-turnover-preventable-often-ignored.aspx.

5. European Commission, "Regulation (EU) 2024/1689 (Artificial Intelligence Act)," Official Journal of the European Union, August 2024. Annex IV specifies the minimum content of technical documentation for high-risk AI systems, including general system description, design specifications, data governance practices, risk management systems, testing and validation procedures, lifecycle change records, and post-market monitoring plans.

6. IBM Security and Ponemon Institute, "Cost of a Data Breach Report 2025," IBM, 2025. Sixty-three percent of organizations reported having no AI governance policies to manage AI deployment or prevent unauthorized AI usage. See https://www.ibm.com/reports/data-breach.

7. Artificial Intelligence Act, Article 11: Technical Documentation. Technical documentation for high-risk AI systems must be prepared before market placement, kept up to date throughout the system's lifecycle, and retained for at least ten years (per Article 18). Annex IV defines the mandatory content. See https://artificialintelligenceact.eu/article/11/ and https://artificialintelligenceact.eu/annex/4/.

8. National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, January 26, 2023. The Govern function establishes policies and accountability structures; the Map function contextualizes AI risks including documentation of system dependencies. Both functions are designed to operate across the full AI lifecycle. See https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf.

9. International Organization for Standardization, "ISO/IEC 42001:2023 — Information Technology — Artificial Intelligence — Management System," 2023. The standard specifies requirements for establishing, implementing, maintaining, and continually improving an AI management system,

including documented policies, processes, and controls for AI system lifecycle management. See https://www.iso.org/standard/81230.html.