

AI Digital Evidence Brief

5 Questions Leaders Should Ask Before Trusting AI-Driven Digital Evidence

A leadership brief on authority, evidentiary basis,
reconstructability, change sensitivity, and accountability.

MICHVI LLP

AI Digital Evidence Brief

5 Questions Leaders Should Ask Before Trusting AI-Driven Digital Evidence

As AI systems, AI agents, and model-assisted workflows increasingly contribute to enterprise decisions and the evidence that supports them, a distinct governance question emerges — separate from model accuracy or technical performance: can the organisation still explain, after the fact, what the AI system contributed, what it relied upon, and whether the resulting evidence can still be reconstructed and evaluated months later? This brief sets out five questions leadership can use to examine that question directly. It is not an AI benchmark, an AI compliance checklist, or a statement about any specific model's performance.

How to Use This

Use this brief as a starting point for a structured conversation among CIO, CISO, Chief AI Officer, CDO, or Internal Audit leadership before AI-generated evidence is treated as a settled basis for a material decision. It does not score or evaluate any specific AI system, model, or vendor.

The Five Questions

1. When an AI system, model, or agent contributes to a decision or generates evidence the organisation relies upon, is it clear what authority that system had — and where a human was required to review, approve, or override the outcome?

CONTRIBUTION AND AUTHORITY

2. Can the organisation identify what evidence, data, or context the AI system relied upon to produce a given output, or does that basis become opaque once the output is used downstream?

EVIDENTIARY BASIS

3. If a decision or piece of evidence produced with AI assistance were challenged later, could the organisation reconstruct the model, prompt, data, tools, and permissions that were active at the time?

RECONSTRUCTABILITY

4. When the underlying model or model version, prompt, connected tools, enterprise data sources, or permissions change, is that treated as a governance-relevant event that may alter the basis of prior AI-assisted decisions — or does it pass unnoticed?

SENSITIVITY TO CHANGE

5. Can the organisation reliably distinguish AI-generated or AI-transformed evidence from the underlying source records or source evidence when both may appear in the same record, report, or decision trail — and is accountability for that distinction clearly assigned?

DISTINGUISHING AI-GENERATED FROM SOURCE EVIDENCE

Closing Interpretation

These questions are not about whether an AI system performs well. They are about whether the organisation can still explain, defend, and reconstruct what that system contributed once the moment has passed. An honest “no” or “unclear” answer to any of these questions is a governance observation, not a technical failure.

Boundary & Disclaimer

This resource is intended for general governance discussion and does not constitute legal advice, compliance certification, audit assurance, or a formal Michvi assessment. It is not a claim about the performance, accuracy, or compliance status of any AI system, model, or vendor.

If these questions surface uncertainty about how AI-driven evidence would hold up under later scrutiny, initiate a confidential governance dialogue.