MortarCloud

A WHITE PAPER / NOVEMBER 2026

The Inception Void

When AI initiatives start in the middle

By Dale Martin

Founder, MortarCloud.ai



Where trust begins

In the final years of my enterprise career, I was responsible for program management and responsible AI governance. Many AI proposals crossed my desk, and I began to recognize a recurring problem.

The project had often started in the middle.

A model had been selected. A data platform was under construction. A pilot was already underway. The next decision was whether to continue investing, expand the deployment, or authorize broader use.

Yet important questions remained unresolved. What business outcome were we trying to improve? Who had authority to decide what the AI system could do? What evidence would demonstrate that the investment was creating value without introducing unacceptable risk?

The business often knew more than the proposal revealed. Its knowledge lived in experienced employees, informal practices, exceptions, and conversations. But that understanding had not been made explicit enough to guide the system or support the decision.

A more capable model could improve the technology. It could not resolve those questions for the business.

I call this missing foundation the Inception Void.

The problem is not confined to organizations beginning their AI journey. It can exist in a mature AI program whenever purpose, knowledge, authority, culture, and evidence are insufficient for the next consequential decision.

Closing the void does not require documenting the entire enterprise before making progress. It requires establishing enough trustworthy business context to make the decision at hand, then retaining and improving that knowledge as the organization and its AI capabilities evolve.

The objective is reliable, bounded, accountable business outcomes. Reliability means performance good enough for the particular purpose, with failure recognized and handled. Boundaries concern the information a system may use, the decisions it may influence, and the actions it may take. Accountability requires a person who can explain the decision to use it, examine what happened, and intervene. An outcome must also respect the people affected by it, including their privacy and safety. Profit alone is an inadequate test of trust.

These expectations require more than selecting a capable model. They require a business to say what matters, what its information means, how authority works, and what conditions would make an outcome acceptable.

The starting line that gets skipped

A leader usually starts where the technology is. I do not say that as a criticism. The technology is where the vendors are, where the demonstrations are, and where the budget line already sits. A reasonable person goes where the path is paved. I also watched disappointing results prompt another search for a better tool when important questions about the work remained unanswered. Better technology could help; it could not settle those questions for the business.

Building MortarCloud.ai has made this principle tangible in my own work. I use AI as an operating system for developing the business, combining intensive context engineering with multiple frontier models. The critical work is not simply choosing the most capable model. It is establishing what the models need to know, which information is authoritative, what decisions remain mine, and how the results are reviewed. When that context is incomplete, the work suffers. When it is explicit and maintained, the results improve. Models and providers may change. The durable asset is the accumulated business knowledge, the decisions behind it, and the human judgment that gives it direction.

The comparison I find useful is onboarding. A new colleague needs to learn what the company is trying to accomplish, who decides, when to ask, and what people mean when they say a job was done well. A database and a login supply only part of that understanding. AI systems also need relevant context, although supplying it does not give them human judgment or make them accountable colleagues.

Nor is this a question that established guidance has overlooked. NIST's AI Risk Management Framework explicitly addresses organizational mission, goals, and business value. Its MAP function informs the initial decision whether to develop or deploy a system and supports reassessment of systems already in use.[1] The gap I am describing occurs when that work has not been done adequately in the business, even though guidance tells us it should be.

Even for a system already in use, the next investment or delegation requires a current answer to those questions.

What the system needs to know

Consider a hypothetical AI assistant that prepares client proposals. It can write well and retrieve the requested documents. That still leaves business questions: which services can be promised, whose approval is required, whether the price list is current, and whether a past exception may be offered again.

An experienced employee might know that an old price was a one-time concession and ask before offering it again. An AI system given the same document can produce an apparently complete proposal without resolving that distinction. The business has to decide which judgments it is delegating, which information can support them, and which decisions still belong to a person. Otherwise, part of its operating practice is being decided by whatever the system happens to infer.

The relevant system includes more than the model. It includes the purpose of the work, the source data, the context supplied with it, prompts, connected tools, agents, the human interface, delegated authority, and the eventual outcome. A change in any of those may change what it is reasonable to trust.

Data fitness belongs inside that assessment. A signed contract may be an accurate record of one agreement without being an approved template for the next. A document can be current yet unavailable to the person receiving the answer. A field called “customer” may mean something different to sales and finance. Permission to retrieve a record does not establish that it is suitable evidence for a particular decision.

Context engineering, in this business sense, means deciding what the system should be given, what that information means, where it came from, and which limitations accompany it. Putting more documents within reach can make a mistaken interpretation easier to repeat if those distinctions remain unresolved.

Culture matters here because the written rule and the working rule may differ. Leaders may describe a formal approval process while staff rely on an informal exception that keeps the work moving. That disagreement needs attention before either account becomes an instruction to an agent. People doing the work must be able to correct the record. Their contributions need a clear route to review, without treating every contribution as an approved fact or a new grant of authority.

Purpose and context also need controls. For the proposal assistant, a practical design might permit drafting while requiring an authorized person to approve a price exception. Access restrictions and approval checks can be enforced by deterministic software where feasible and tested at the point of action. The model's instruction to behave responsibly is insufficient by itself.

Those controls can constrain a probabilistic component without making its outputs deterministic or guaranteeing every business outcome. Likewise, a fluent explanation from a model is not proof that its answer is correct; generated explanations and citations can themselves be false.[2] Accountability should rest on records and checks we can actually inspect, including inputs, permissions, tool use, decisions, and results. It should not depend on a promise of complete access to opaque model reasoning.

What the business should measure

A vendor's claim can be accurate and still leave the business question unanswered. A model may summarize documents well. An agent may route tickets correctly. A platform may reduce the time taken for a task under the conditions of its test. A stack of true sentences about tools does not yet tell a leader whether this business will be better off, by how much, or how anyone will know.

I have sat in the meeting where a leader tried to connect a vendor's slide with a quarter's numbers and lacked a reliable way to do it. The leader was not careless. The evidence needed to connect the investment to the business result had not been planned with the same care as the investment itself.

That is why the financial baseline remains central to the Inception Void. Before changing a process, record what it costs, how long it takes, how often it goes wrong, and what quality is acceptable. Connect those measures to a business purpose. Faster proposal drafting is useful only in relation to something the business values, such as better response, more productive capacity, or sounder commercial decisions.

The measures must include the work introduced by AI. Drafting time can fall while review and correction rise. Time released from a task creates capacity; it becomes a cash saving only when spending actually changes. A useful assessment considers the costs of implementation, operation, oversight, and failure, as well as the benefit being claimed.

A dated baseline does not promise that the AI will pay back. It gives the business a reference point. The costs of errors, correction, review, and cybersecurity exposure can change the value case after launch; the same decision record should make both expected benefit and accepted risk visible. Establishing cause requires more: a credible account of what would otherwise have happened, attention to changes in workload or staffing, and a comparison suited to the decision. A simple improvement after deployment does not establish that AI caused it.

If AI is already in use, the opportunity to evaluate it has not disappeared. Historical records may support a partial reconstruction. An appropriate comparison group or bounded test may help. A current baseline can support evaluation of the next change, provided it is not mislabeled as the original pre-AI state. Government guidance on evaluating AI interventions asks evaluators to assess what baseline data can be collected retrospectively and to prioritize evidence appropriate to the evaluation.[3]

Be clear about what the evidence can and cannot establish. Sometimes the responsible conclusion will be that a result is promising but its cause is uncertain. That is a better basis for a decision than a precise return calculation built on assumptions nobody has named.

Gartner has announced a November 4, 2026 webinar examining cost, cybersecurity, trust, business context, human oversight, and governance as AI scales. That published agenda is relevant context for this problem, not a finding from a completed webinar or an endorsement of MortarCloud.ai.[8]

A business record that survives the technology

An organization can answer these questions for one project and still lose the answers when people leave or a platform changes. Business sovereignty requires a usable record of purpose, knowledge, decisions, and performance that remains under the organization's control. That record need not be finished all at once. A bounded decision can become the first reviewed contribution, followed by other useful contributions as the business learns what it needs.

The business should define what its AI systems and supporting software need to do. Its accumulated understanding should help it choose and change technology, and should remain available when a particular tool reaches the end of its usefulness.

This is broader than asking what happens if a vendor disappears. Ordinary changes matter: a better model becomes available, a contract ends, a business unit is sold, or an experienced employee leaves. The practical question is what the organization retains and can continue to use. Does it still have the meanings behind its data, the reasons for decisions, the sources supporting its knowledge, and the authority to maintain the working assets?

A file export is only part of an answer. Moving a table does not necessarily preserve its relationships, history, permissions, or the behavior of the application that uses it. Portability needs a practical test: can a competent replacement team understand the record, run the delivered system, maintain it, and identify which dependencies must be replaced? What time and cost does that require?

The same standard should apply to a consultant. An engagement should leave the business with an asset it can use and maintain, with the rights and documentation necessary to do so. Proprietary methods can support how the consultant creates that asset. Continued use of the completed client work should not require an indefinite relationship with its creator.

For MortarCloud.ai, this is a delivery requirement: when a client asks us to stop, we will stop. The client must be able to retain the working software and business record created for it and operate them independently of MortarCloud.ai. Future services and proprietary production tools are a separate matter. That requirement must be supported by the delivered architecture, continuing rights, documented dependencies, and a tested handover. A tested handover is the acceptance test for that commitment.[4]

This is a standard that must be met for any engagement that includes working software. Our current documented exit package preserves client-owned records and static views, but has not demonstrated that an independent operator can install and rebuild the proprietary toolchain without our assistance. The software rights, dependencies, and handover therefore must be agreed and tested before a claim of operational independence is made for a particular client.

No format alone establishes durability. A successful handover provides evidence under the conditions tested; changes in the business and its technology can require the test to be repeated.

Closing the void when you have already started

Most executives reading this will have AI in use. Step Zero should help them make a sound next decision without requiring a restart of everything they have built.

Step Zero is the work of establishing sufficient purpose, context, baseline, data fitness, accountability, and cultural alignment for the decision at hand. It can begin with an existing use case. The necessary depth depends on the consequences of the work, and it expands when dependencies or risks demand it. The business does not have to finish describing itself before improving one bounded activity. Step Zero is an enduring discipline of establishing or verifying enough context to decide whether an AI capability should exist, what it may do, who may authorize its actions, and how evidence of value and harm will be reviewed. It is not a mandatory consulting stage or a demand to restart a working AI program.

For the proposal assistant, that could mean clarifying the commercial objective, identifying authoritative material, establishing approval rights, and measuring the work as it actually happens. Those answers can become part of a wider business record, rather than disappearing into another project's documents.

Someone also needs responsibility for keeping those answers usable. A new price, a changed approval limit, or a disputed source can require review before the next use. The work is to identify changes that matter to the decision, resolve competing accounts through the appropriate business owner, and record what was accepted and when. Capturing more information without that responsibility can simply create another record to distrust. Reviewed entries should remain dated, attributable, correctable, and useful for the next decision; as more entries are accepted, the client-owned foundation can grow without forcing the whole business to be mapped first.

MortarCloud.ai organizes its advisory work around five connected service capabilities: Ground, Align, Plan, Build, and Keep. The Trusted Living Agent Framework™ is our governing method, not the name of the commercial service model. Ground establishes the business context and what is already in use. Align relates that activity to purpose and identifies gaps. Plan sets priorities, responsibilities, and controls. Build brings execution evidence from the client's team or chosen builders into the record. Keep supports review and maintenance as conditions change.[5]

We treat those capabilities as flexible points of entry and return, determined by the business need, while the scope and acceptance criteria of each engagement are agreed separately. A live use case may expose missing context that requires Ground; new evidence may send an implementation back to Align or Plan. This preserves the five capabilities while avoiding a requirement that every engagement traverse them in the same order.

Existing work then becomes something to assess on purpose. A useful agent may continue within appropriate limits. Another may need a narrower role, better information, or a pause. Serving a stated goal is one consideration; fitness, authority, and acceptable consequences also matter.

Making the business record visible

The picture I would want as a leader connects goals, people, processes, data, systems, agents, and vendors, with sources behind the connections and unresolved gaps visible. A knowledge graph provides a way to record those relationships. Its value comes from the decisions it helps the business make and the quality of the record it preserves.

For the proposal assistant, a useful record would let an owner trace a price to its source, the person authorized to approve an exception, and the uses that depend on it. When the policy changes, those connections give the business a basis for deciding what needs review. A concise record may be enough for one bounded task. A richer graph earns its place when shared dependencies make those decisions easier to make and maintain. Its upkeep belongs in the value calculation too.

MortarCloud.ai's Trusted Living Twin™ makes that record visible in its internal demonstration. Client-environment deployment is an intended design; independent client installation and continued operation have not yet been verified. The business should be able to see what is known, what needs attention, and

how a proposed change relates to the work it is trying to accomplish. An incomplete but honestly sourced picture can be useful. It should never be presented as a complete inventory merely because its visualization looks finished.

Our working knowledge-building process has four parts. AI-assisted extraction reads selected documents and exports. People review the proposed entries, or enter information directly with a source and date. A deterministic build produces the versioned graph from those reviewed inputs. The Twin's pages render from that graph. With the reviewed inputs and build version held fixed, the build is reproducible; that does not make the source information correct.[6] Selected and reviewed portions of the record could be connected progressively to a client's enterprise AI environment as suitable interfaces, permissions, and operating checks are implemented and tested; that integration is not established by the present demonstration.

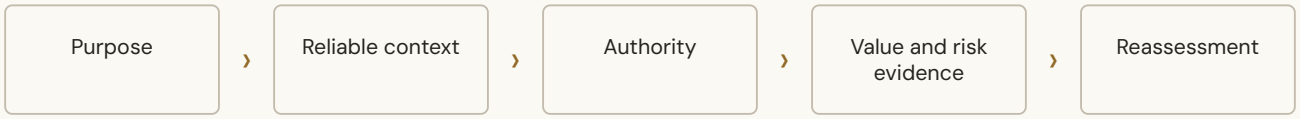
Live connectors, continuous automated monitoring, universal discovery of AI in the organization, and enforcement at the moment an agent acts are planned capabilities. They should not be inferred from a working graph and its pages. A record can describe a control without implementing it, and a recent source can describe a process that no longer works as intended.

Hollister Reed LLP is a fictional accounting firm used to demonstrate the same knowledge-building process with simulated financial and operational records. A fictional firm lets us show that detail without exposing a real business's confidential information. The demonstration illustrates the process; it does not report client results.[6]

That distinction belongs in the record itself. A model's proposal, a person's acceptance, a link to a source, and an independent verification are different things. They can coexist without meaning the same thing. Making those differences visible gives a leader a basis for deciding what to rely on and what still needs checking.

Five questions to ask this week

Five questions for one consequential AI decision



A decision aid, not a claim of implemented software controls.

These questions can reveal where to begin, whether or not you engage a consultant:

1. What is this part of the business trying to accomplish, and would the people directing its AI describe success in compatible terms?

2. For one important AI use case, which information is authoritative and fit for the decision, what may the system do with it, and who can correct the record or stop the activity?
3. Which tools, agents, and vendors participate in that work, what authority have they been given, and where are ownership or approval still unresolved?
4. What evidence connects the work to business value, including quality, cost, risk, and the people affected, and what does our comparison leave uncertain?
5. If we changed a major model, vendor, or adviser, what business knowledge and working assets would remain usable, and when did we last test that?

Write the answers down, date them, and assign the unresolved questions. That begins the work; it does not certify the system. The useful result is a clearer decision about what to do next, supported by a business record that can be corrected and carried forward.

If this applies to you

Ground is MortarCloud.ai's bounded engagement for establishing context, a baseline, and a record of the AI already in use. It is one possible starting point for the work described here. The aim is to leave the business with a usable record and a clearer basis for its own decisions, whether the engagement continues or ends. Further essays are at dalemartin.ai; information about MortarCloud.ai is at MortarCloud.ai.

Sources and notes

1. National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, January 2023, section 5.2 and MAP 1.3–1.4, pp. 24–26. Primary source. The framework supports the statement about existing guidance; the observed gap in practice is the author's professional observation.
2. National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1, July 2024, section 2.2, "Confabulation." Primary source. This source addresses false generated content, including logic and citations; it does not validate any MortarCloud.ai implementation.
3. UK government, *Guidance on the Impact Evaluation of AI Interventions*, published with the Magenta Book, HTML edition updated May 15, 2026, sections 2.1–2.4, especially 2.3, "Establishing a baseline." Primary source. This is public-sector evaluation guidance; the measurement recommendations here are the author's application to business decisions.
4. The client exit requirement reflects the founder's direction of October 30, 2026. It is a delivery commitment to be implemented and tested, not an independently demonstrated portability result. Specific delivered rights and dependencies must support the commitment.
5. Ground, Align, Plan, Build, and Keep are the established service names in MortarCloud.ai's October 2026 canonical records. The flexible hub-and-spoke model and enduring Step Zero discipline reflect founder-approved November 2 strategy decisions; their implementation scope and results remain subject to evidence and engagement acceptance. Build retains its current scope of execution evidence from the client's team or chosen builders.
6. MortarCloud.ai, October 30, 2026 strategy handoff, Part 07, "Capability and Evidence Register," read with the Twin architecture and manuscript source notes. These records document the working extraction, human-review, graph-build, and rendering process. Hollister Reed LLP uses simulated records. The complete evidence and provenance ledger accompanies this manuscript.
7. This manuscript develops *The Inception Void*, draft 0.1, and the approved draft 0.2. The account of AI as an operating system for business development, using intensive context engineering and multiple frontier models, describes the author's firsthand operating practice; it is not a measured comparison of providers or a claim of client results. The original manuscripts and review material are preserved.
8. Gartner, *The AI Value Gap Is Really an AI Risk Gap: Protecting Value at Scale*, webinar announcement for November 4, 2026 (forthcoming when this manuscript was prepared on November 2). Publisher announcement. The announced agenda provides independently published issue framing only. It does not imply Gartner endorsement, completed findings, or validation of MortarCloud.ai.

About the author

Dale Martin is the founder of MortarCloud.ai and author of *Bears & Mosquitoes: Building Trusted AI from the Inside Out*.

He spent three decades in enterprise technology, including senior management responsibility for AI operations and Responsible AI governance at Thermo Fisher Scientific. His responsibilities spanned internally developed generative AI capabilities, the transition to an enterprise ChatGPT Enterprise architecture, and governance within a global organization of more than 120,000 colleagues.

At MortarCloud.ai, he applies that experience to business-first AI, context engineering, and client-owned business knowledge.

The views expressed are his own and do not imply endorsement by any current or former employer or affiliated organization.

Contact: dale@dalemartin.ai | MortarCloud.ai