MortarCloud

A WHITE PAPER / NOVEMBER 2026

The Business Comes First

Putting AI to work and building knowledge the business can keep

By Dale Martin

Founder, MortarCloud.ai

The executive question

A team has an AI assistant, a custom application and several ways to share information. Work still stalls. A service promise is unclear, an old exception appears in a new proposal, or nobody knows who can approve a change. People reach for another tool to fill the gap. The business gains capability while its understanding of how the work fits together falls behind.

At MortarCloud.ai, we help leaders resolve those gaps around a consequential business decision. We connect relevant knowledge, people, systems and authority, establish what can be relied on, and leave a working record the business can use again. Our proprietary intake methods already support bounded collection and review of selected business records and contributions. We are developing broader federated participation and automation so client-owned knowledge can evolve incrementally while accountable owners validate it.

A practical starting point is Ground, our bounded AI Value & Trust Baseline. For an agreed workflow or decision, it establishes the relevant context, known AI activity, responsibilities, evidence and unresolved questions. Its first result is a clearer decision and reusable client-owned business knowledge. Clients already operating AI may instead enter through Align, Plan, Build or Keep, with a proportionate Step Zero check before action. Where an operational change requires implementation, the client or chosen builder executes it against an agreed scope and acceptance criteria.

The longer-term opportunity is cumulative. Each useful increment can leave behind knowledge that helps with the next decision, the next application and the next change of tool. The business should gain useful capability while retaining the understanding and assets on which that capability depends.

Three connected concepts organize our approach. Bears & Mosquitoes establishes the overall problem. The Inception Void addresses the foundation for building. Non-Human Resources describes the shift toward a co-intelligent workforce. The Trusted Living Agent Framework™ provides the central governing approach; our intake develops reviewed knowledge, and the Trusted Living Twin™ makes that record usable.

The question for a leader is practical: where can we make useful progress now, what must we understand to do it well, and what will the business retain?

Step Zero is an enduring business discipline, not a sixth required purchase. Before choosing or changing AI, we establish or verify sufficient understanding of the mission, objectives, culture, operating reality, constraints and authority relevant to the decision. That accepted context becomes part of the client's retained record and can be revisited as conditions change. A first Ground engagement can begin with one consequential decision, normally involving one or two departments and a bounded source set. The first brick should resolve a useful question while laying a foundation that the business can extend.

1. Bears & Mosquitoes establishes the problem

Bears & Mosquitoes examines a change in the activity businesses must govern. Established approaches focus on visible systems, suppliers and boundaries. AI also places capability inside ordinary work, where numerous interactions and delegated actions can have consequences beyond the individual task. The same distributed capability creates opportunity: people close to a problem can investigate it, develop an idea and improve a service.[1]

The management task is to connect that activity to business purpose, dependable information and appropriate authority. Approved tools can be used with incomplete context. A shared assistant can become a dependency without a maintenance owner. An accurate search result can retrieve an old agreement without establishing whether its terms apply to a new commitment. Existing security and access controls remain necessary; business context helps explain what happens within and across them.

Technology sprawl becomes a business problem when overlapping tools, unclear responsibilities or disconnected information make work harder to understand and maintain. Shadow AI adds uses outside the organization's appropriate visibility or approval. Neither problem is explained adequately by counting applications. A specialist tool may serve a justified need; an unreviewed use of an approved product may still create a consequential gap.

A workaround is therefore a useful place to investigate. Does the employee lack a capability, access to current knowledge, a workable handoff, or awareness of an existing function? Is the approved route too slow to resolve an exception? Understanding the reason gives the business a basis for improving the arrangement.

The response might be to reuse an existing capability, change a process, configure or integrate a tool, approve a specialist service, commission a bounded build, or retire something redundant. The relevant record should preserve why the choice was made, who owns it and what would justify reconsidering it.

Our current intake can examine selected records and contributions within an agreed scope. The resulting account must make its coverage and unknowns clear. Technical discovery of undisclosed tools and continuous monitoring require their own capabilities and evidence. A well-organized record should never imply visibility beyond its sources.

For the business, the benefit to pursue is a more responsive, accountable way to address gaps. People need a practical route from an observed problem to a decision and an appropriate response.

2. The Inception Void establishes the foundation for building

The Inception Void names the gap between the context an AI system needs and the context the business has made explicit, usable and current. Useful context includes objectives, terminology, relationships, policies, responsibilities, authority and the meaning of information. It also includes the current way of working and a dated baseline against which a change can be assessed.

Building includes buying and configuring a product, connecting systems, adopting an assistant and changing an existing process. Each choice shapes how work operates. The foundation helps the business decide what should change and what the result must accomplish.

This work can begin around an existing deployment. A firm does not need a complete model of itself before improving one activity. It needs sufficient context to resolve the consequential questions for that activity, including dependencies that could change the decision. Unknowns that prevent responsible progress should lead to a narrower scope, further evidence or a reasoned stop.

NIST's AI Risk Management Framework connects AI assessment with organizational mission, goals and business value, and supports reassessment as conditions change.[2] Context engineering adds a useful technical and organizational perspective: how relevant information and the conditions for its use are prepared and maintained for an AI task. Gartner's public article emphasizes accountability for context, validation, governance and continuing updates, including the governance of knowledge graphs.[3]

The business question behind those terms is who can establish what information means, judge whether it applies here, and correct it when circumstances change. A policy document, a historical transaction and an employee's explanation can each contribute something different. Their agreement cannot be assumed.

Our own operating practice makes this concrete. ChatGPT is a core component of how we develop MortarCloud.ai and has been in use for some time. We have also used Claude. Preparing reusable business context has helped us adapt our tool choices relatively quickly in that work. This is firsthand operating experience. A client's change still requires examination of its integrations, permissions, instructions, dependencies and accepted results.[4]

The foundation serves decisions before and during implementation. It can inform an existing enterprise architecture practice or support a firm working without a formal architecture program. Its immediate purpose is to make the next investment or change more specific, accountable and useful.

3. Non-Human Resources describes a co-intelligent workforce

Non-Human Resources describes a broader change in how human capability and AI capability participate in business work. It joins two responsibilities: managing AI capabilities through their working lives, and designing the human roles, learning, coordination and supervision around them. Every consequential delegation to AI is also a decision about human work.

Ethan Mollick's work on co-intelligence and management provides an important reference for human direction, expertise, delegation and evaluation. That intellectual contribution is his. We apply our governance and business-knowledge approach to the organizational questions that follow.[5]

An employee may know how to create a useful assistant without having authority to connect it to confidential records. A manager may ask for more output while review capacity becomes the constraint. A senior colleague may leave while an assistant continues to depend on the judgment that colleague supplied. Naming an owner in a register does not resolve these conditions. Ownership requires knowledge, time, evidence and authority to intervene.

Knowledge building belongs in this work design. Leaders establish priorities and constraints. People close to the work identify missing information, practical exceptions and improvement opportunities. Relevant owners resolve contested meanings and accept changes. A useful contribution needs a visible response: accepted, corrected, deferred or declined with a reason.

The intended employee benefit is practical. A coordinator should have a clearer route to resolve a recurring ambiguity. A reviewer should be able to inspect the accepted interpretation and its source. A junior colleague should be able to learn why an exception requires escalation. These benefits depend on making contribution and review part of supported work, with suitable tools and time.

The organizational implications extend beyond a single workflow. Shared AI capabilities need maintenance, while each consequential use retains a business owner. People need a professional home for development even when working across functions. Leaders must allocate review capacity, resolve priorities and decide how additional capacity will be used.

Greater output is only one possible benefit. Better service, learning, quality and new opportunities also matter. If AI removes tasks through which employees developed judgment, the learning path needs attention. Asking people to contribute what they know should support their ability to do useful work and participate in its improvement. Headcount reduction should not become the organizing purpose.

4. Our proprietary federated knowledge intake system

Our proprietary federated knowledge intake system turns distributed business knowledge into a connected, sourced record. It combines a working knowledge-building process with additional automation and broader participation that we are actively developing. It is also a knowledge intake and onboarding system: a way to establish the context needed for a defined piece of work.[4]

Knowledge enters from the people and records relevant to that question. Documents describe policies and services. Selected system records and exports show activity. People explain what those records mean, where exceptions matter and how work actually happens. The process attends to source, scope, ownership and authority as well as content.

Today we use AI-assisted extraction from selected documents and exports, human review, and direct sourced entry. AI can propose connections and draft entries. People add context, examine conflicts and decide what belongs in the accepted business record. A responsible owner resolves a business interpretation or leaves the uncertainty explicit. Reviewed inputs produce a versioned record through a repeatable build, and the Trusted Living Twin™ presents usable views.[4]

Proprietary describes our methods and tools. Federated describes distributed contribution and responsibility: relevant knowledge belongs across the business, including the front line. Automation assists the process and is being expanded. These attributes do not imply universal connectors, automatic reconciliation, cross-client pooling or a completed employee contribution platform.

A source link shows where a statement came from. Human acceptance records a decision about its use. An independent check supplies a different kind of evidence. Keeping those distinctions visible helps the business challenge and improve the record. Repeating a build consistently does not establish that every accepted interpretation is correct.

The resulting knowledge belongs to the client under the agreed engagement terms. Its value is in supporting a bounded decision and preserving understanding that can be maintained and reused. Our proprietary production methods remain distinct from that client asset.

Federated intake contributes to the broader discipline of context engineering. Supplying appropriate context during live AI execution additionally requires selection, delivery, access controls, integration and maintenance in the actual operating environment. Those functions must be implemented and tested for the intended use; they do not follow automatically from constructing a reviewed record.

5. A fictional line of sight from a proposal gap to useful work

The following professional-services firm and sequence are fictional. They illustrate how the current knowledge-building process could support a bounded engagement, followed by possible implementation and reuse. They are not a customer case study or a report of measured results. The later workflow changes would require separate implementation and testing; this example does not describe additional features already demonstrated by our software.

The problem the firm brings

The firm uses an assistant to help draft client proposals. A proposal coordinator repeatedly has to ask whether a standard service includes a particular integration. The current service description excludes it. An older signed proposal includes it. A colleague's local template repeats the older promise.

The coordinator sometimes reaches for another assistant because the existing workflow does not make the relevant answer easy to find. Management initially sees a tool-choice problem. The coordinator's account reveals a recurring uncertainty about the service itself. Experienced reviewers repeatedly resolve the same questions, while the practice lead is considering extending AI-assisted proposal preparation across the service team.

The bounded question is what the firm can delegate in this class of proposal, under what service boundary, and with whose approval and review capacity. That decision matters before participation expands. The firm does not begin by mapping every department or replacing its applications.

The knowledge that changes the decision

The agreed inputs are the current service description, the earlier proposal, the local template, the relevant approval policy and explanations from the coordinator and practice lead. Permission to use those inputs is settled within the engagement scope. AI-assisted preparation can identify the apparent conflict and propose relationships among the service, past exception and approval responsibility.

In this fictional case, human review establishes that the older integration was a specifically approved exception. It remains evidence of that earlier commitment. It does not authorize the same promise in every new proposal. The practice lead accepts a clear interpretation of the current service boundary and identifies who can decide on a new exception.

The record retains the sources, the accepted meaning, its owner and the conditions requiring further review. If delivery capacity for a requested exception remains unknown, that question stays open. It does not disappear because the draft sounds convincing.

A result the firm can use before another build

In this fictional illustration, a bounded Ground engagement produces a reviewed account of the service boundary and decision authority, documented dependencies, and a dated baseline using the available evidence. A static Twin-style view can make those reviewed relationships easier to inspect where included in the agreed deliverables; independently deployed client Twin software remains subject to separate technical and contractual acceptance. Missing measures remain visible.

The practice lead can now decide what drafting may continue within the standard scope, when an exception needs approval and what review capacity is required before participation expands. The proposal process owner is assigned to correct or retire the misleading local template. The coordinator has an accepted interpretation to consult. A new reviewer can see the reasoning and source, rather than reconstructing the earlier conversation. Those are useful outputs even if the firm commissions no integration.

Illustrative retained knowledge	Decision it supports
Current service definition and source	What the proposal may describe as standard
Earlier exception and its limited scope	What cannot be assumed to apply again
Accountable owner and approval responsibility	Who can resolve a new exception
Open delivery question and dated baseline	What must be learned before extending the work

The possible implementation

The firm might then ask its chosen builder to configure the existing proposal workspace around the accepted sources, display the relevant exception guidance and retain the required approval step. The implementation brief would specify the permitted behavior and the evidence needed to accept it. Tests could include an ordinary proposal, a requested exception and an unresolved service question.

That technical change is a separate piece of work. Its permissions, failure handling and controls must work in the actual application. A policy recorded in the Twin does not itself prevent an assistant from making an unauthorized promise.

The firm would evaluate the combined result: preparation, review, correction and approval effort; accepted quality; and the full cost of the arrangement. Faster drafting would be useful only if the complete process serves the intended outcome.

The next use and the next contribution

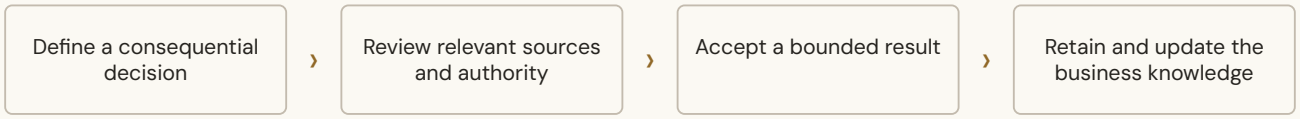
For a later proposal, the coordinator could reuse the accepted service definition and approval responsibilities where they still apply. The next question starts from that reviewed account. A different client circumstance might require new knowledge; a policy change might require a correction.

The same service boundary could later inform onboarding, with its applicability checked by the onboarding owner. This is the intended accumulation of value: useful knowledge becomes available to related work while new differences remain visible.

The coordinator's original contribution also receives a response that improves the work. Continuing participation can identify the next gap. Whether this saves time, improves quality or sustains participation is something the firm would measure, not an outcome assumed by the illustration.

6. The asset that connects successive improvements

How a useful business decision extends the foundation



A conceptual progression, not a deployed technical workflow or an additional service model.

A connected business record can link goals, people, processes, data, systems, agents, vendors, decisions and evidence. A knowledge graph is one way to organize those relationships. An ontology supplies agreed concepts and relationships; a semantic layer makes their business meaning usable. Their value depends on the questions they help people answer.

This foundation can grow through a hub-and-spoke approach. The hub is shared, reviewed business knowledge. Each spoke is a bounded application of that knowledge to valuable work. A spoke may improve an existing process, use a purchased capability or involve custom implementation. The next use extends the shared record only where its needs justify the effort.

The hub is a description of how knowledge and decisions connect. It does not require every application to route through one technical gateway or every routine decision to wait for central approval. Local owners need authority to act within agreed limits, with a way to resolve consequential conflicts across the business.

Our architecture keeps four parts distinct:

Part	Contribution
Proprietary federated knowledge intake system and methods/tools	Develop reviewed knowledge through the working process; further automation and participation are being developed
Client-owned organizational knowledge	Preserve accepted meanings, relationships, sources, decisions and open questions
Trusted Living Twin™	Make the reviewed record usable through software views intended for the client environment

Part	Contribution
Trusted Living Agent Framework™	Govern the approach through purpose, context, authority, accountability and evidence

The Framework remains our central governing IP. Its five pillars are Identity, Data, Infrastructure, Action and Orchestration. Its signatures are Self-describing, Self-reporting and Self-governing. These concepts guide examination and design; a named pillar or signature does not establish an implemented control.

Several technical distinctions matter. GraphRAG is a retrieval approach that uses graph-based information to help assemble context for an AI response. A business knowledge graph can support ownership and decision review without operating GraphRAG.[6] MCP can connect an AI system to tools and information, but appropriate authorization, security and governance remain necessary.[7] Predefined workflows and model-directed agents also require different attention to delegated discretion and handoffs.[8]

We currently build and render the reviewed record. Verified live Twin context delivery, runtime authorization, action blocking, universal agent discovery and continuous automated monitoring are not established capabilities of that process. The Hollister Reed LLP demonstration uses fictional, simulated records and includes an agent register; it is not evidence of a deployed client inventory.[4]

Gartner's public service-delivery abstract argues for embedding context engineering in AI services. Its public context-platform abstract describes platforms that construct, govern and deliver reusable context to agents during execution.[9][10] These sources support the relevance of the problem. They do not establish our implementation, uniqueness or comparative advantage. Our contribution must be assessed through the usefulness of our approach and outputs, including client ownership, human judgment and total delivery effort.

7. Make knowledge building an ongoing business practice

A first engagement should establish useful business context that remains intelligible and editable after the project. Its continuing value depends on how the business uses and maintains it. We intend to develop this context into a governed, client-controlled knowledge foundation, enlarged brick by brick rather than through a compulsory company-wide census.

A practical cycle begins when ordinary work exposes a gap: an employee finds an exception, a reviewer challenges an interpretation, a service changes or an owner leaves. The relevant people contribute the information, an accountable owner reviews the change, and the accepted record is updated for affected uses. NIST's AI RMF supports clear responsibilities, regular review and the incorporation of adjudicated feedback from relevant participants.[2]

This is the ongoing practice we are developing around our intake system. A limited first engagement may involve a department lead and selected contributors; later federation should allow the appropriate parts of the workforce to propose changes without silently making each submission authoritative. Broader participation and further automation should make contribution and review easier, while preserving the distinction between a suggestion, supporting evidence, accepted knowledge and permission to act.

Maintenance needs a boundary. The business should know which knowledge is being maintained, who notices changes, what triggers review and what happens when freshness is uncertain. A polished view of an obsolete rule can mislead. More entries and more submissions are useful only when they improve decisions enough to justify their review and upkeep.

The business also needs to decide how much knowledge to share across activities. Common concepts can be reused; meaningful differences between teams should remain visible. An enterprise-wide definition is not required before every local improvement. The people who use and maintain the knowledge should help determine the appropriate scope.

Durability includes the ability to continue through changes in people, tools and advisers. Today, the documented delivery package can preserve the agreed business record, selected source snapshots and mappings, and static copies of accepted views. Independently rebuildable software remains a proposed deliverable subject to a documented client license, installed-package test and non-founder exit rehearsal. A maintained document or spreadsheet may be the better retained asset when it answers the decision without additional complexity.[4]

A meaningful handover test asks a qualified replacement operator to use the record, correct a business fact and maintain the accepted functions without MortarCloud.ai service access. Viewing an old record is only part of that requirement. During consulting, any licensed software use should be contractually specified. Afterward, a client may choose an optional continuation service or retain and extend its accepted knowledge through its own personnel or another adviser. We will not charge merely for access to the client's accepted business knowledge; independently usable proprietary software remains subject to demonstrated capability and explicit rights.

8. Begin with a result worth buying

Ground, Align, Plan, Build and Keep are connected capabilities. They may overlap, and a firm already using AI can begin with the gap in its current work. The starting scope should fit a decision that matters now.

Capability	Business question
Ground	What do we know about this work, and what evidence is missing?
Align	What outcome matters, who owns it and what authority applies?
Plan	Which option is justified, and what would make it acceptable?
Build	What evidence shows that the client or chosen builder implemented the agreed change?
Keep	What must remain current, who maintains it and when should it be reviewed?

Ground provides a bounded way to begin through the AI Value & Trust Baseline. Before work begins, we agree the business question, relevant workflow and source boundary, client participation, outputs, timing, fee and acceptance responsibility. An initial engagement usually centers on one important decision

involving one or two departments, subject to actual complexity. Its scope can address an existing problem without requiring a company-wide architecture exercise or five consecutive purchases. Broader work should be scoped separately rather than presumed to fit a ten-department package.

The result should give the sponsor four things: a reviewed account of the relevant work and knowledge; a dated baseline and visible gaps; a usable business record with accountable owners; and a justified next decision with requirements for any implementation. The sponsor should be able to explain what that result enables and why it is useful even if no further service is purchased.

An operational change may need a client team, vendor or other builder. Responsibility, acceptance, support and cost for that implementation must be explicit. We examine execution evidence within the agreed Build scope. A broader implementation obligation requires its own agreement.

We judge value and risk together through the combined work. A potentially valuable AI system can still fail because of unsuitable context, unclear authority, weak security, uncontrolled cost or inadequate oversight; conversely, appropriate governance helps preserve value as use expands. This is consistent with the agenda Gartner announced for its November 4, 2026 webinar, which is external context rather than evidence validating our method.[11] Speed means time to an accepted useful result, including preparation, review and correction. Quality means fitness for the business purpose and handling of consequential errors. Durability means that knowledge and working assets remain understandable, maintainable and reusable through change. Adaptability means that a tool can be changed with the necessary adjustments and checks while useful business context is retained.

A dated baseline helps evaluate those benefits. It does not establish cause on its own: workload, case mix, staffing and other process changes may affect the result. Time saved creates capacity before it creates cash savings; the business needs a plan to use that capacity. Client effort, our effort, tool costs, correction and ongoing maintenance all belong in the assessment.

The next related use provides a further test. What knowledge was reused? What had to change? Did the retained record reduce repeated clarification or support a better decision? If a simpler approach serves the work adequately, that is a valid outcome. Continuing services should earn their place through identifiable work and continuing value.

Trust grows from useful results, clear limits, accountable people and evidence that can be examined and corrected. Those conditions need to remain visible as capability expands.

9. Bring one consequential piece of work

A useful first conversation starts with something specific: proposals that need repeated correction, an assistant whose authority is unclear, a handoff that depends on one person's memory, or a tool decision whose business requirements remain unsettled.

Bring the outcome you want, the people responsible for the work, the sources available and the point at which uncertainty prevents progress. We can then define whether a bounded Ground engagement would produce a result worth buying, what participation it would require and what should happen next.

Bears & Mosquitoes explains the overall problem. The Inception Void establishes the foundation for building. Non-Human Resources describes the co-intelligent workforce paradigm. MortarCloud.ai connects them through a current knowledge-building process, a governing Framework and a client-owned working asset that can develop with the business.

The aim is useful progress on the work in front of us, with knowledge the business can carry into what comes next.

Make sense of AI change. Stay in control.

Understand your business. See where AI fits. Keep the knowledge and the decisions yours.

Sources and notes

1. Dale Martin, *Bears & Mosquitoes: Building Trusted AI from the Inside Out*, book website and Built for the Wrong Threat. Source for the conceptual distinction between visible systems and distributed AI activity, including its business opportunity. The metaphor is not a measured risk classification. Older numerical illustrations and stronger capability claims are not used.
2. National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, January 2023. Official publication. MAP 1.3-1.4 addresses mission, goals and value; GOVERN 1.5-1.7, 2.1 and 5.1-5.2 address review, inventories, retirement, responsibilities and feedback. The business application here is the author's interpretation.
3. Avivah Litan, *Context Engineering Is the New Prompt Engineering*, Gartner, October 6, 2025. Public article. Page metadata: *Context engineering: Why it's Replacing Prompt Engineering for Enterprise AI Success*. Used for organizational accountability, validation, governance and maintenance of context. These recommendations do not establish MortarCloud.ai outcomes.
4. MortarCloud.ai, October 2026 working-process and architecture records; The Inception Void, content 0.4; Non-Human Resources, approved content 0.2 and designed edition 0.3; Dale Martin's subsequent founder directions. Company passages distinguish documented practice, founder experience, development work and delivery requirements. ChatGPT use and relative ease of adapting tools are founder-reported experience, not a measured migration comparison. The extended proposal sequence is fictional and separate from the existing simulated Hollister Reed LLP demonstration. Independent client operation remains a requirement to implement and test. Private source identities are retained in the review register.
5. Ethan Mollick, *Management as AI superpower*, January 27, 2026. Author's essay. Reference for direction, expertise, delegation and evaluation. Co-intelligence is attributed to Mollick's contribution and is not claimed as MortarCloud.ai IP.
6. Microsoft, GraphRAG documentation and Query Engine overview. Public documentation reviewed for the retrieval distinction. It does not establish a MortarCloud.ai GraphRAG deployment.
7. Model Context Protocol project, Authorization and Security Best Practices, versioned 2026-07-28 guidance. Used for authorization boundaries and least privilege; applicable controls depend on the implementation and transport. Protocol use alone does not establish compliance or business authority.
8. Anthropic, *Building effective agents*, December 19, 2024. Primary technical guidance. Used for the distinction between predefined workflows and model-directed agents, including their composition. Later tooling updates are not relied on here.
9. Ben Fieselmann and Tom Coshow, *AI Services Race: Win by Embedding Context Engineering in Delivery*, Gartner, June 16, 2026. Public abstract. Only the public summary was reviewed; the full report and supporting research were not accessed.
10. Kjell Carlsson, Afraz Jaffri, Christopher Long and Radu Miclaus, *Market Overview for AI Context Platforms*, Gartner, August 13, 2026. Public abstract. Only the public summary was reviewed. Its description of reusable context during execution is not a claim about MortarCloud.ai's present runtime capabilities. No Gartner endorsement, matrix placement or adoption forecast is asserted.
11. Gartner, *The AI Value Gap Is Really an AI Risk Gap: Protecting Value at Scale*, announced webinar scheduled November 4, 2026. Public webinar description. The announcement identifies cost, cybersecurity, trust, business context and continuous oversight as issues; this paper does not claim access to the future presentation or Gartner endorsement.

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