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. We work with professional-services firms of roughly one hundred to one thousand people, where the knowledge that matters sits in documents, past engagements and the heads of experienced people.

The 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. A firm already operating AI can enter through Align, Plan, Build or Keep, with a proportionate Step Zero check before action. Where an operational change requires implementation, the client or its chosen builder executes it against an agreed scope and acceptance criteria.

The longer-term opportunity is cumulative. Each useful increment leaves behind knowledge that helps with the next decision, the next application and the next change of tool. The business gains 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; federated knowledge intake develops reviewed knowledge across the business, 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 begins with one consequential decision, normally involving one or two departments and a bounded source set. The first brick resolves 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 explains 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 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 preserves why the choice was made, who owns it and what would justify reconsidering it.

Our intake examines selected records and contributions within an agreed scope, and the resulting account states its coverage and its unknowns. A well-organized record never implies 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 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 on the subject 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 each contribute something different. Their agreement cannot be assumed.

Our own operating practice makes this concrete. We run MortarCloud.ai on written business context that both ChatGPT and Claude read and write, and section 5 describes how. 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 has a clearer route to resolve a recurring ambiguity. A reviewer can inspect the accepted interpretation and its source. A junior colleague can 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. Federated knowledge intake

Federated knowledge intake turns distributed business knowledge into a connected, sourced record that the business keeps feeding. Federated is the operative word, and it is the part of our method that does the most work. The knowledge of a business is already spread across it: in the service lead's head, the coordinator's template, the finance team's exceptions, the front line's workarounds, the old proposal that still gets copied. Intake does not pull that knowledge into a central project and hand it back. It spreads the work of contributing and reviewing across the people who hold the knowledge, as part of the work they already do, so that the record is fed continuously rather than assembled once.[4]

Why central knowledge projects fail

Every firm of any age has tried to write itself down. The intranet, the wiki, the process library, the knowledge base: each was launched with a project team, filled in a burst, and left to age. The pattern fails for three reasons, and none of them is a lack of effort.

The central team cannot know the exceptions. The people who know why the standard service excludes an integration, or why one client's terms differ, are the people doing the work, and they were interviewed once. The knowledge ages from the day it is written. A service changes, a policy is revised, an owner leaves, and the entry stays as it was, polished and wrong. And contribution is extra work. Writing to the knowledge base sits outside the job, so it happens last, and the person who does it gets nothing back.

AI has made the failure more expensive. An assistant pointed at that library answers confidently from the stale entry, and the business is bound by the answer. The Air Canada case in section 5 is the small, public version of a problem every firm with an assistant now has.

How the federated loop works

Federation turns the central project into a loop that runs inside ordinary work.

A gap surfaces in the work. A coordinator cannot tell which service description governs. A reviewer challenges an interpretation. A new hire asks why an exception needs escalation. A senior colleague announces a departure. Each is the moment the knowledge is live and the person who holds it is present.

The person closest to the gap contributes. They say what they know, point at the source, and name the uncertainty. AI drafts the entry, links the sources, and proposes how the new fact relates to what the record already holds. The contributor confirms or corrects the draft. This takes minutes, not an afternoon, because the AI does the writing and the person does the knowing.

An accountable owner reviews. Every part of the record has an owner with the knowledge, time, evidence and authority to decide. The owner accepts the contribution, corrects it, defers it with a date, or declines it with a reason. The contributor sees the response. A contribution that disappears into a queue is the last contribution that person makes; a contribution that comes back accepted, or corrected with the reason, is the start of a habit.

The record updates for every use that depends on it. The reviewed entry is versioned, dated, and attributed to its contributor and its owner. The views the business uses, in the Trusted Living Twin™ or in the documents the firm already works from, carry the accepted version. The next person with the same question finds the answer, its source, and who accepted it.

The next gap starts the loop again. The record is never finished, and it is never a project. It is maintained the way the firm maintains its client list, its rate card and its service catalogue: by the people whose work depends on it, as part of that work.

Embedded in the work, not added to it

The difference between federation and the knowledge base is where the contribution happens. Federation puts it at the moment of work: when an exception is resolved, when a proposal is corrected, when a new colleague is onboarded, when a tool is changed. Those moments already exist in every firm. Federation gives each one a place to leave what was learned.

Leaders set the priorities and the boundaries: which decisions matter now, which parts of the business the record covers first, who owns what. People close to the work identify what is missing, where exceptions matter, and how the work actually happens. Owners resolve contested meanings and accept changes. A light stewarding function keeps the loop running: it notices stale entries, routes contributions to the right owner, and reports what the record covers and what it does not.

Knowledge enters from the people and records relevant to a 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. We extract from selected documents and exports with AI assistance, review by hand, and enter sourced facts directly. AI proposes connections and drafts

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.[4]

What the record becomes

Fed this way, the record becomes a second brain for the business: a living account of what the firm knows, who knows it, where it came from, and what is still open. It is a knowledge engine in the plain sense that it is fed continuously and produces answers continuously, and it belongs to the business that feeds it.

Federation is what makes it durable. No single team carries the load, so no single departure empties it. Each part of the business feeds the part of the record it knows best, and the record says who contributed each part and who accepted it. When a senior colleague leaves, the judgment they supplied is in the record with their name on it, and the owner who accepted it is still there.

It is also what makes it honest. 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 lets the business challenge and improve the record. A suggestion, supporting evidence, accepted knowledge and permission to act are four different things, and the record keeps them apart. An accepted interpretation is a decision, and decisions can be revisited.

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 production methods remain distinct from that client asset.

Federated intake contributes to the broader discipline of context engineering. Supplying context during live AI execution additionally requires selection, delivery, access controls, integration and maintenance in the actual operating environment, and those functions are implemented and tested in the operating environment for the intended use.

5. Three things that happened, and how we work

The gap this paper describes is not hypothetical. Three public events show it from three angles, and our own operation shows the alternative.

The assistant that made a promise the business had not authorized

In 2022 a passenger asked Air Canada's website chatbot about bereavement fares. The chatbot told him he could buy a full-fare ticket and claim the discount afterward. The airline's own policy page said the opposite. When he claimed, the airline refused, and in February 2024 the Civil Resolution Tribunal of British Columbia ordered it to pay. The airline had argued that the chatbot was a separate legal entity responsible for its own statements. The tribunal rejected that argument and found that the airline had not taken reasonable care to make its chatbot accurate.[9]

The sum was small. The lesson is the one in section 1: the business had the right answer on one page and the wrong answer in the assistant, and nobody owned the difference. A policy recorded anywhere does not by itself prevent an assistant from making an unauthorized promise. Someone has to establish which statement governs, connect the assistant to it, and own the result.

The deliverable where review capacity was the constraint

In July 2025 Deloitte Australia delivered a review of a welfare compliance system to the Australian Department of Employment and Workplace Relations under a contract worth about A$440,000. A university academic found that the report cited academic works that did not exist and attributed a quotation to a Federal Court judgment that did not contain it. Deloitte reissued the report with an appendix disclosing that a generative AI model had been used in its preparation, stated that the corrections did not change its findings, and agreed to repay the final installment of the contract.[10]

This is the condition section 3 names: more output was asked for, and review capacity became the constraint. The firm had the expertise to catch every error. The work was not designed so that the expertise was applied before the client saw the result.

The delegation decided on cost and reversed on quality

In early 2024 Klarna announced that its AI assistant was doing the work of about seven hundred customer-service agents. In May 2025 its chief executive told Bloomberg that organizing support around cost had led to lower quality, and that the company was hiring human agents again while keeping AI for routine volume.[11]

Time saved creates capacity before it creates cash savings. Klarna counted the savings first and the quality second, and then paid to put the quality back. The order matters.

The pattern that works

Morgan Stanley's wealth management business took the opposite route when it introduced an assistant for its advisors in 2023. It bounded the assistant to the firm's own library of research and procedures, about one hundred thousand documents, and built evaluations that tested the assistant's answers against real advisor questions before each use was released. The firm and OpenAI have described the method publicly: a scoped corpus, evaluation before release, and expert feedback at each step.[12]

That is the shape of the intake in section 4, built by a firm with the resources to staff a central team for it. Most firms do not have those resources, and federation is the answer: the same bounded question, the same reviewed sources, the same owner who accepts the result, with the work of contributing and reviewing spread across the people who already hold the knowledge.

How we work

MortarCloud.ai runs on its own method. The firm's strategy, decisions, contacts, interactions, drafts and publication masters are written records in version control. ChatGPT and Claude both read those records and both write to them. A handoff from one tool to the other is a file, not a conversation. When one tool does a piece of work better than the other, we change the tool, because the business context does not live in the tool.

The method began in 2022 and has been the firm's operating system since the start of 2026. It produced this paper and the two that follow it. Our public demonstration, Hollister Reed LLP, applies the same method to a fictional firm's simulated records so that a visitor can see the record, the register and the views without any client's data.[4]

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

6. The asset that connects successive improvements

A connected business record links 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 grows 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
Federated knowledge intake	Develop reviewed knowledge across the business as part of ordinary work
Client-owned organizational knowledge	Preserve accepted meanings, relationships, sources, decisions and open questions
Trusted Living Twin™	Present the reviewed record as usable views; a client license for the software is optional and agreed after an engagement
Trusted Living Agent Framework™	Govern the approach through purpose, context, authority, accountability and evidence

The Framework is 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; the controls they call for are implemented in each engagement and in each client's environment.

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 supports ownership and decision review without operating GraphRAG.[6] MCP connects an AI system to tools and information, and authorization, security and governance remain the business's responsibility.[7] Predefined workflows and model-directed agents require different attention to delegated discretion and handoffs.[8]

We build and render the reviewed record. Our contribution is assessed through the usefulness of that record and the decisions it supports, including client ownership, human judgment and total delivery effort.

7. Make knowledge building an ongoing business practice

A first engagement establishes useful business context that remains intelligible and editable after the project. Its continuing value depends on how the business uses and maintains it. We develop that 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]

A first engagement involves a department lead and selected contributors. Federation grows from there: each part of the workforce proposes changes to the part of the record it knows, and no submission silently becomes authoritative. The distinction between a suggestion, supporting evidence, accepted knowledge and permission to act stays visible throughout.

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 misleads. 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 help determine the appropriate scope.

Durability includes the ability to continue through changes in people, tools and advisers. The delivery package preserves the agreed business record, selected source snapshots and mappings, and static copies of accepted views. A maintained document or spreadsheet may be the better retained asset when it answers the decision without additional complexity.[4]

The handover test is simple: a qualified replacement operator uses the record, corrects a business fact and maintains the accepted functions without MortarCloud.ai. Afterward, a client may choose an optional continuation service or extend its accepted knowledge through its own people or another adviser. We do not charge for access to the client's own accepted business knowledge.

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 fits 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 agreed change was implemented?
Keep	What must remain current, who maintains it and when should it be reviewed?

Ground is the 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 centers on one important decision involving one or two departments. Its scope addresses an existing problem without a company-wide architecture exercise or five consecutive purchases. Broader work is scoped separately.

The result gives 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 can 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 are made 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; appropriate governance preserves value as use expands. 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. Workload, case mix, staffing and other process changes also affect the result, and the baseline records them. 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 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 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 federated knowledge intake, a governing Framework and a client-owned working asset that develops 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.

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 distinction between visible systems and distributed AI activity, including its business opportunity.
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 is the author's.
3. Avivah Litan, *Context Engineering Is the New Prompt Engineering*, Gartner, October 6, 2025. Public article. Source for organizational accountability, validation, governance and maintenance of context.
4. MortarCloud.ai working-process and architecture records, 2022 to 2026; The Inception Void, publication master v1.1; Non-Human Resources, publication master v1.1. The firm's own records are the source for sections 4, 5 and 7. The Hollister Reed LLP demonstration uses simulated records of a fictional firm.
5. Ethan Mollick, *Management as AI superpower*, January 27, 2026. Author's essay. Reference for direction, expertise, delegation and evaluation. Co-intelligence is Mollick's contribution.
6. Microsoft, GraphRAG documentation and Query Engine overview. Source for the retrieval distinction.
7. Model Context Protocol project, Authorization and Security Best Practices, versioned 2026-07-28. Source for authorization boundaries and least privilege.
8. Anthropic, *Building effective agents*, December 19, 2024. Primary technical guidance. Source for the distinction between predefined workflows and model-directed agents.
9. *Moffatt v. Air Canada*, 2024 BCCRT 149, Civil Resolution Tribunal of British Columbia, February 14, 2024. Summaries: Manatt, Mondaq, CBS News.
10. Associated Press, *Deloitte to partially refund Australian government for report with apparent AI-generated errors*, October 7, 2025, via BNN Bloomberg; Fast Company; Cyber Daily. The contract value is as reported by the Department of Employment and Workplace Relations.
11. Bloomberg interview with Sebastian Siemiatkowski, May 2025, as reported by Entrepreneur, Computing and Maginative. The seven-hundred-agent figure is Klarna's own, from its February 2024 announcement.
12. OpenAI, Morgan Stanley customer story; CNBC, Morgan Stanley is testing an OpenAI-powered chatbot for its 16,000 financial advisors, March 14, 2023. The figure of about 100,000 documents is CNBC's, from the firm's internal memo, and is repeated by Morgan Stanley's head of firmwide AI product and architecture strategy in OpenAI's account.

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