MortarCloud

A WHITE PAPER / NOVEMBER 2026

Non-Human Resources

The workforce your firm never hired, and the function that manages it

By Dale Martin

Founder, MortarCloud.ai

1. When the work changes

In a hypothetical finance team, an invoice discrepancy appears just before a payment is due. An artificial intelligence (AI) assistant compares the records and prepares a recommendation. A workflow routes the case for approval, while an agent chooses how to investigate a disputed item within the limits set by the business. The reviewer sees a plausible answer. Can that person inspect the evidence, challenge the agent's choices, and stop the payment if necessary? The system may be able to do more of the work. That does not mean the organization has decided who should carry the responsibility.

The software supplier may offer more autonomy, but consequential actions can still require human approval. The human work has already changed. Someone must decide which records can be relied on, which exceptions require judgment, and what the system may do before a person looks. The colleague who once assembled the case may now review a conclusion. The manager must decide whether that colleague has the time, knowledge, and authority to challenge it. If the builder leaves, someone else must understand enough of the work to take responsibility for it.

The organization chart may look much as it did before. The responsibilities beneath it have changed.

I use AI as an operating system for developing MortarCloud.ai. That practice relies on intensive context engineering: keeping relevant business knowledge, decisions, sources, and constraints usable across the work. I use multiple frontier models interchangeably, selecting the one suited to a task rather than building the business around any provider. When business context is missing, the deficiencies become apparent. Doing the inception work first improves the quality of my work and raises a larger management question: once AI is involved, how should responsibility and authority be organized?

That question reaches beyond the agent that can act on its own. It includes the employee who builds a useful assistant, the operational team that comes to depend on it, and the leader who uses AI to examine a decision. Each arrangement can expand what people are able to do. Each also changes what they must understand, what they are responsible for, and where they need support.

Non-Human Resources is the management function that brings these responsibilities together. One part manages AI capabilities through their working life: purpose, selection, introduction, performance, limits, change, and retirement. The other redesigns human work around those capabilities: roles, skills, learning, collaboration, management, and the terms on which people are asked to rely on AI. These two parts need to be reviewed together whenever the work or its authority changes.

Trust is the North Star. In the finance example, the test is whether discrepancies are resolved accurately and in time, whether the right people can question an exception, and whether useful capacity is created after review and correction are counted. The people doing the work also need to know how their roles will change, how their work will be assessed, and how they will develop the judgment the new arrangement requires.

Every material delegation to AI is also a decision about human work.

2. One function with two responsibilities

The capability and the people around it should enter the same management conversation. Buying a system, naming an owner, and publishing a policy do not settle who will investigate its errors, maintain the relevant expertise, or absorb the additional review work. Those responsibilities need resources and authority as well as names.

Capability management asks whether the AI is suitable for its purpose, what it may access and do, how it is evaluated, and when it must change or stop. Human work design asks who contributes, how tasks and decisions move between people and systems, what people must learn, and how their work will be recognized. A permission change may require a new reviewer. A staffing change may make the previous level of delegation untenable. An apparently successful automation may be removing the practice through which new colleagues learn.

This function need not become a new department. Business leadership owns the purpose and outcome. The people function and line managers address roles, capacity, learning, and fair treatment. Technology, architecture, and security establish feasible systems and controls. Data owners and domain experts establish meaning and suitability. Risk and governance help challenge consequential uses and examine evidence. Their work must converge on a clear decision about a particular use.

These are related but different responsibilities. People remain responsible for setting objectives, making judgments beyond the delegated boundary, caring for employees, assigning resources, and authorizing consequential changes. AI capabilities may perform defined tasks or exercise bounded discretion through implemented systems; they do not acquire employment status, a career manager, or independent corporate accountability. Organizational accountability remains with the business and its authorized people, including when a supplier operates the underlying system. Calling a capability a teammate or digital worker can help discussion, but it must not change those decision rights.

Leadership should assign an existing executive responsibility for making capability management and human work design operate together across workflows. That executive must be able to secure decisions and resources from the contributing functions; individual workflow owners remain responsible for applying the practice to their work. A small firm may combine these responsibilities in one person. A larger firm must keep the coordination clear even when several functions contribute.

The Inception Void, the companion paper, concerns the business context required for trustworthy AI outcomes. Non-Human Resources concerns how people and AI systems work within that context. Purpose and knowledge inform the work; the experience of doing the work reveals what must be corrected, learned, or redesigned.

Step Zero gives this work a business and cultural starting point. Before assigning or expanding an AI responsibility, the organization establishes enough of its mission, objectives, values, existing practices, decision rights, sources, constraints, and affected people to judge that particular use. Step Zero is revisited as the work changes; it is not a compulsory transformation program or a new department. An organization already operating agents can begin with a consequential existing delegation, examine the

missing context, and correct that bounded arrangement rather than restart its entire AI program. The point is to design the delegation for the business, not redesign the business merely to accommodate a tool.[13]

3. People become builders and different kinds of leaders

Ethan Mollick's work on co-intelligence and delegation offers a useful starting point: effective work with AI depends on human direction, subject knowledge, and evaluation.[1] A co-intelligent organization makes that relationship part of its ordinary work. The cyborg metaphor captures close collaboration; it does not transfer the organization's accountability to a system.

An operational employee may use AI to explain a discrepancy, compare alternatives, or build a small application. A manager may explore scenarios before allocating resources. A professional may ask AI to critique an argument while retaining responsibility for the advice. These arrangements can change the practical boundary of a role before its job description changes.

A one-day product-development experiment at Procter & Gamble found that individuals using generative AI produced proposals of comparable assessed quality to those of two-person teams working without it. That finding invites organizations to reconsider how expertise is combined for a task. It does not establish that permanent teams, specialist knowledge, or management can be removed.[2]

Citizen development makes the distinction between creating capability and authorizing its use especially important. An employee may be able to build an invoice assistant without having authority to connect it to shared financial records or send its conclusions to a customer. As the work becomes shared or consequential, the organization should provide a clear route to review, support, and approval. Otherwise, it asks the builder to carry responsibilities that the job may never have included.

Executives belong inside this change. A CEO should be able to frame a consequential question, work with an approved AI system, inspect its evidence, and recognize when specialist judgment is needed. An assistant can help coordinate the work, but cannot supply the executive's understanding or accountability on the executive's behalf.

Administrative professionals may also redesign their work toward coordination, exception handling, institutional knowledge, and judgment. Assuming that their roles are universally obsolete would conceal the actual design question: which contributions remain valuable, what can be delegated, and what additional responsibility deserves training and recognition?

The leadership requirement is participation with judgment. Leaders who expect others to change should demonstrate the competence and working practices they require. That expectation will need different practical standards for a chief executive, an analyst, a service employee, and a specialist reviewer.

4. Organize around the work

What should the organization look like when people can draw on AI capabilities across traditional functional boundaries? Adding agent boxes below existing managers is an incomplete answer. Reporting lines identify only part of the arrangement. They do not show which workflow may invoke a shared capability, whose data it uses, what authority applies, or who can stop a consequential action.

I propose beginning with one bounded workflow inside the existing organization. Give its continuing business outcome a named owner, supported by professional homes, shared AI and data services, and explicit leadership and assurance authority. Assign these responsibilities to existing people before considering new reporting lines. This tests a way of organizing work across functions; it does not presume a universal successor to hierarchy or matrix organization.

Organizational design already has a substantial history of balancing more self-contained work with shared expertise and lateral coordination. Galbraith's information-processing view makes those tradeoffs explicit. Research on less-hierarchical organizations likewise treats decentralization as a deliberate redistribution of authority, with limits that require attention.[3] AI gives us another reason to revisit the choices; it does not remove them.

A clear home for people and an owner for the work

Outcome teams take responsibility for a continuing business result, such as resolving invoice exceptions accurately and promptly. Each has a named human workflow owner who coordinates the contributions of people and AI. Membership can change as the work requires, while ownership of the result remains clear. A temporary project can use the same principle with a defined end and handover.

Professional homes develop people and maintain expertise. Each person has one clear career manager responsible for development, feedback, staffing, and succession. A home might be finance, customer operations, engineering, or another professional community. It commits agreed capacity and learning time to the work; it does not issue a competing daily list of priorities without resolving the allocation.

Shared AI, data, and platform services maintain reusable capabilities and their service boundaries. A shared invoice-analysis capability may support several outcome teams. Its service owner is responsible for the capability's operation and material changes. Each consuming workflow still needs an accountable business owner for its particular use. Sharing a capability does not merge every user's authority into one permission.

Leadership retains decisions about priorities, investment, acceptable risk, and contested resources. Assurance has defined review and intervention responsibilities proportionate to the consequences. In a smaller firm, existing people can hold several roles, provided workload and the need for independent challenge are addressed. This proposal does not require a new layer of executives or committees.

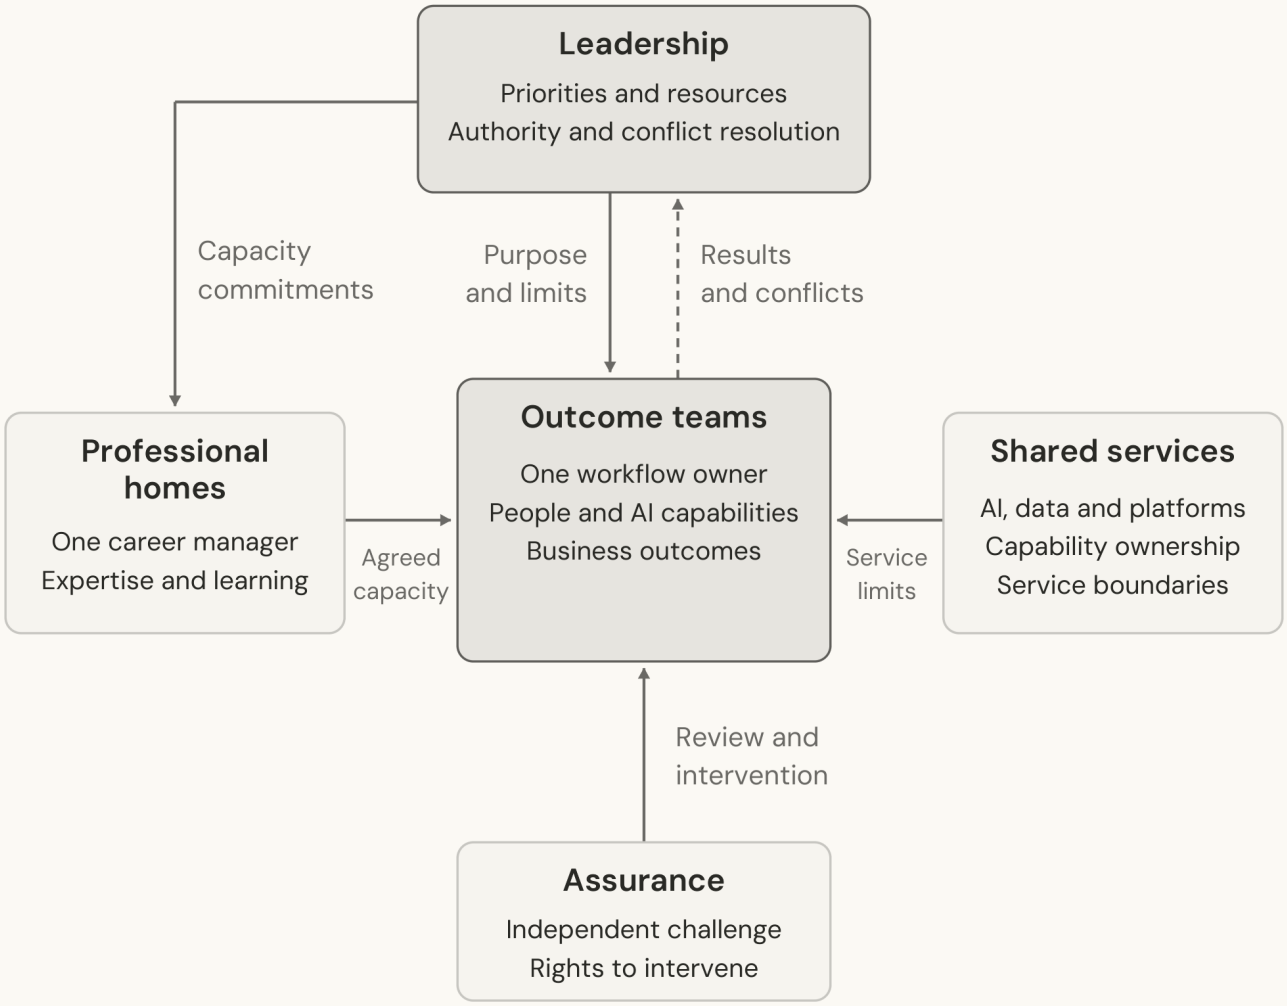


Figure 1. A proposed operating structure for human-AI work. People retain a clear professional home; business outcomes have named owners. AI capabilities appear in the work and delegation view. This is an explanatory proposal, not a validated organizational template or a new MortarCloud.ai reference architecture.

Decision rights make the structure real

The practical distinction is who can make which decision. Within an agreed allocation, the workflow owner sets delivery priorities. The career manager owns development and career decisions. A named leader resolves demands that exceed the agreed allocation. Employees should not have to reconcile incompatible instructions by working longer hours.

The workflow owner may authorize a business use within delegated limits. Technical and data owners implement the corresponding access and service boundaries. A reviewer may withhold approval or invoke an agreed intervention. If the proposed use exceeds an accepted service or data boundary, the workflow owner must obtain a decision from the person authorized to change that boundary. A contested exception

goes to the named operating or risk authority. Until it is resolved, the existing restriction or agreed safe fallback applies. Sharing a capability cannot expand another workflow's permissions, and a productivity target cannot silently override a reviewer's intervention.

This arrangement retains functional and outcome dimensions. It can become a matrix with new labels if both dimensions issue competing commands, or if shared services become an unaccountable queue. Its value must come from clearer work allocation, faster resolution of conflicts, maintained expertise, and evidence of better outcomes. A stable functional structure may remain appropriate where the work is predictable and well contained. The proposed change should earn its place in the business.

In the hypothetical finance workflow, a finance professional retains a development home and commits time to the invoice-resolution team. The workflow owner decides which discrepancies qualify for delegated handling. A shared-service owner maintains the AI capability. A reviewer handles exceptions with authority to pause its use. Leadership resolves a conflict if the new review burden exceeds available capacity. The organization can then see a management problem that an agent inventory alone would miss.

AI capabilities should not be counted as people in the reporting hierarchy. A reusable capability can serve multiple workflows; an execution can be temporary; a supplier's internal agents may be opaque. The organization needs a comprehensible view of accountable humans, shared capabilities, and consequential delegations. It does not need to invent an employee relationship for every software execution.

5. Authority must follow the work

Workflow and agent describe different aspects of the arrangement. A workflow coordinates steps toward a result; those steps may use rules, probabilistic AI, or agents. An agentic component can select how to pursue an objective, including which steps or tools to use, while remaining subject to supervision and limits. Agents can also invoke workflows. Anthropic's technical guidance distinguishes paths set in code from model-directed activity and also describes combined patterns.[10] These are useful working distinctions, rather than a complete risk classification.

Routing an invoice under established rules and choosing how to investigate an exception can sit in the same process. Using a language model alone does not establish a delegated digital role. Meaningful discretion over operational choices creates an additional need to define that role and its authority. The business still needs an owner and appropriate controls for the whole process, including any consequential rules-based action.

Before authorizing either arrangement, identify which choices are delegated, which require a person, what evidence will be kept, and who can suspend the use. Where an agent and a workflow hand work to each other, the receiving step must stay within the permitted purpose, access, and actions of that handoff. An agent should not gain broader authority merely by calling another system. The delegator cannot grant authority they do not hold. The amount of autonomy is a business decision about delegated authority.

Delegation also depends on the business knowledge supporting the task. Reviewed, source-linked accounts of a process, its exceptions, owners, policies, and current constraints can help people define suitable role boundaries and provide relevant context to AI systems. That account must distinguish an AI-

proposed assertion, a cited source, a business owner's acceptance, and a separately tested operating condition. Knowledge can inform a permission decision; it cannot itself grant access, prevent an action, or enforce an escalation. In a federated contribution approach, people closest to the work can identify gaps, while authorized owners resolve contested interpretations before the accepted record is used more widely.[13]

The first authority question is what the system may cause to happen. Summarizing an invoice, recommending a correction, changing a record, and releasing a payment are different delegations. Their boundaries should follow the purpose and consequences of the action. A long period without a reported error does not, by itself, justify broader authority.

Gartner's public guidance uses data sensitivity and autonomy as starting considerations for agent oversight. Its broader guidance addresses permissions, scope, ownership, meaningful approval, and organizational accountability.[4] These considerations belong in a business decision about the work. They are not evidence that every highly autonomous use has greater residual risk than every advisory use.

Instructions, permissions, and controls must remain distinguishable. A prompt can describe a limit. A permission determines access. An implemented control may prevent an action. Evidence from operation helps establish whether the intended limit holds. Writing a boundary in a policy or business record does not show that a system enforces it.

The same distinction applies to visibility. A business may control a capability it builds, coordinate temporary agents through an orchestration system, or consume a vendor service whose internals it cannot inspect. It should record the boundary honestly. It can govern its own access, integrations, contracts, and decision to continue using the service without claiming control over every internal agent.

The aviation idea of a Tail Number offers a useful analogy for identity. For the arrangement proposed here, I recommend applying it to a persistent capability, while recording executions and grants of authority separately. The record should connect the work to the responsible people and the relevant evidence. This is a proposed management distinction; its technical implementation depends on the system being used.

Changes need the same discipline as introductions. A new model version, data source, tool, supplier behavior, or permission can change the arrangement. The responsible people should assess what that change means for the work before extending reliance. Retirement must address access, dependencies, retained records, and who takes over any remaining obligations.

6. Oversight people can actually perform

A named human is not sufficient oversight. The person needs relevant knowledge, usable evidence, enough time, and practical authority to question or interrupt the work. Assigning responsibility without those conditions can conceal a management failure behind an ownership field.

Trustworthy reliance grows when the person can test a conclusion against usable evidence and intervene before the consequences exceed their authority. Delegated autonomy should therefore be adjusted to both the action's consequences and the real review capacity available, not simply to the system's apparent

fluency. Reassessment is needed when models, data, permissions, staff, suppliers, or business objectives change. A reviewer who can only acknowledge a polished result, without enough knowledge or time to challenge it, is not providing meaningful oversight.

The distinction matters even when the system only advises. A reviewer who sees a polished conclusion may still need to reconstruct its evidence. A manager who expands the number of cases processed must account for the exception and review work created. If the business cannot provide the necessary capacity, it should narrow the delegation, improve the arrangement, or stop the use.

Gartner's autonomy guidance explicitly warns that approvals can weaken under time pressure and fatigue.[4] The response should be appropriate to the task: useful review interfaces, clear escalation, access to original evidence, and feasible interruption or recovery. Universal approval of every step can be impractical; the absence of a human approval step can also be unacceptable for a particular consequence. The organization must decide deliberately.

The employee analogy needs a similar boundary. Familiar disciplines such as role definition, supervision, review, and handover can help organize the discussion. AI cannot take the organization's responsibility for choosing and authorizing its use. A working paper on AI employee framing found small average effects on error detection, with a stronger, exploratory pattern among managers whose organizations already put AI on organization or workflow charts. It did not test whether a named-owner record solves the problem.[5]

Continuity is part of oversight. When a builder, owner, or reviewer leaves, the business needs a successor with the knowledge and capacity to act. Until then, a consequential use may require narrower authority or suspension. A vacant name on a record should not leave a system operating on an assumption that someone is still watching.

The record protects the person as much as the firm. It can show what was authorized, what information was available, what was reviewed, and where a concern was escalated. That supports fair examination of events. It does not establish that all decisions were sound or absolve anyone of responsibilities beyond the record.

7. A culture of contribution

A co-intelligent organization should make role-relevant AI competence a requirement, beginning with its executives. People should learn to use approved capabilities where those capabilities improve their work. Leaders should demonstrate that competence themselves. Participation does not mean accepting every output or forcing AI into every task. Choosing a human approach, challenging a result, or stopping a harmful action can demonstrate better judgment than using more AI.

That expectation requires reciprocal commitments from leadership. The organization should provide suitable tools, access, learning time, examples relevant to the role, and support when the work changes. It should explain what employees are expected to learn, what evidence will demonstrate readiness, and how concerns will be handled. Role requirements should accommodate the actual work and legitimate access or accessibility needs.

The starting purpose should be additional value: better service, stronger capability, improved quality, new possibilities, and more useful capacity. Cost discipline belongs in that conversation. Headcount reduction should not be the organizing purpose or the sole definition of success.

The tension is practical. We ask people to contribute the knowledge that makes AI useful, identify its mistakes, and help redesign their work. If they believe the main reward for doing so is eliminating their own role, we give them a reason to withhold what the organization needs. A growth narrative will have little credibility if the targets, incentives, and decisions consistently tell a different story.

Research supports taking that risk seriously without treating every situation as identical. Studies of teams connect psychological safety with learning and improvement. A separate three-wave study of Korean employees associated perceived job insecurity with lower psychological safety and more knowledge hiding.[6] These findings do not prove that every cost-focused AI program destroys culture. They strengthen the case for honest commitments and for protecting the ability to speak up.

Leadership should state how expected gains will be used. Possibilities include serving unmet demand, improving response times, developing new offerings, removing avoidable work, and giving people time to learn or tackle more difficult problems. Those commitments need to appear in budgets and work allocation. Asking people to learn while keeping every existing demand unchanged is not a credible investment in capability.

Workforce changes may still occur. The organization should consider development, redeployment, and normal attrition where appropriate, while explaining material changes honestly. It should not promise that every job will remain unchanged. Nor should it treat attrition as automatically benign: leaving a role unfilled can transfer workload, remove expertise, or weaken review capacity. The effects on the remaining work need to be assessed.

Make competence visible in ordinary work

An executive can demonstrate readiness by using AI to examine a business question, checking the evidence, identifying the limits, and explaining the decision retained by the human. An operational employee can demonstrate it through a completed case and the reasoning behind an exception. A reviewer can show that they recognize a misleading output and know how to intervene. These are examples of role-relevant assessment, not a new certification scheme.

Managers should assess results and judgment, not prompt counts or declarations of enthusiasm. Knowledge sharing, useful challenges, corrected assumptions, and responsible escalation should count as contribution. A mistake reported early should become an opportunity to examine the work; concealing a known problem should not be rewarded by a clean-looking dashboard.

Learning also needs work of its own. In a randomized study of 52 developers learning an unfamiliar Python library, AI-assisted participants scored lower on an immediate follow-up assessment, while average task-completion time did not improve significantly.[7] The study does not establish inevitable long-term deskilling. It does show why completing a task and learning to understand it must be evaluated separately.

Professional homes therefore need to preserve mentoring, supervised difficult cases, explanation, and opportunities to demonstrate independent judgment. Removing routine work without redesigning the route to expertise can leave the business dependent on systems that fewer people understand. The cultural commitment should be to developing capable people alongside useful AI.

8. Judge the combined work

A faster AI task is one contribution to a business result. The relevant unit of assessment is the whole workflow, including the people and systems required to produce an acceptable outcome.

The same logic links risk to value. Gartner's public October 2026 article treats cost, cybersecurity, and trustworthiness as interconnected obstacles to protecting value at scale. Its announced November 4, 2026 webinar names business context, governance, and human oversight as subjects for discussion; as this manuscript is prepared on November 2, that event has not occurred.[12] This is independent context for the management problem, not evidence that Gartner endorses Non-Human Resources, MortarCloud.ai, or any claimed product result. A faster agent that moves error correction, uncertainty, or review labor elsewhere in the business may not create an improved outcome.

In the finance example, compare the quality and timeliness of resolved discrepancies, the time spent reviewing and correcting them, the handling of exceptions, and the effects on customers and staff. Include software and integration costs, shared-service support, and the effort needed to maintain reliable records. A shorter drafting time is useful only in the context of what happens next.

Research gives reason to expect different outcomes across tasks and workers. A study of 5,172 customer-support workers at one software company estimated a 15% average increase in issues resolved per hour with an AI assistant, with larger benefits for less experienced and lower-performing workers.[8] That is a result for a particular service arrangement, not a forecast for every business or a direct measure of profit.

Establish a baseline and compare the redesigned work with it. A before-and-after comparison does not establish how much of the change AI caused. Account for changes in demand, case mix, staffing, and process; use a credible comparator where feasible. Look for useful capacity, quality, responsiveness, cost, learning, and effects on the people involved. Investigate where apparent gains have moved effort into review, rework, or another team. Make the intended use of released capacity explicit so that an improvement can become value rather than an unexplained gap in the staffing plan.

Trust should remain a condition of success throughout. An output that meets a local performance target while exceeding authority, mishandling information, or preventing meaningful challenge is a reason to change the arrangement. Improved throughput does not compensate for every other failure.

9. Put the function to work

Begin with one important workflow. Identify its purpose, the people affected, the AI capabilities involved, and the decisions that matter. Make the human and AI contributions visible, including the work needed for review and correction. Identify who can authorize, change, question, or stop the use, and provide the resources those responsibilities require.

The workflow may already be running. A client can begin by examining an existing assistant, supplier agent, exception process, or a narrow decision whose AI involvement has expanded beyond its original assumptions. Establish enough Step Zero context to evaluate its current role, authorities, evidence, and effects on people. Keep what is justified, constrain what is not, and document what must be tested or returned to a human. A bounded correction can be useful without a company-wide redesign or a commitment to purchase every service stage.[13]

Then examine the organization around it. Is there one accountable workflow owner? Do people have clear work allocations and a development home? Who maintains shared capabilities? Who resolves competing demands? Can reviewers obtain evidence and intervene? What happens when a responsible person leaves? These questions make organizational design concrete before a company redraws its entire chart.

Agree what better performance would mean and how gains would be used. Assess the combined work, address problems, and expand only when the evidence and human capacity support the next scope. NIST's voluntary AI Risk Management Framework reinforces the need for clear organizational responsibilities, training, and differentiated roles in human-AI arrangements and oversight.[9]

At MortarCloud.ai, the Trusted Living Agent Framework™ connects business purpose and context with five pillars: Identity, Data, Infrastructure, Action, and Orchestration. Non-Human Resources brings the management of AI capabilities and the redesign of human work into that conversation. The service capabilities Ground, Align, Plan, Build, and Keep support the work across its life; the business still owns its decisions and outcomes.[11]

The Trusted Living Twin™ makes reviewed business knowledge and relationships usable. Its documented current work includes selected-file extraction, human review or direct sourced entry, graph construction from reviewed inputs, and presentation. A record can show intended authority and supporting evidence. The current Twin does not provide runtime authorization, action blocking, live discovery, or continuous monitoring. Those functions would require separately implemented and verified systems and controls.[11]

The distinction is deliberate. The Trusted Living Agent Framework™ is the governing method for examining identity, data, infrastructure, action, and orchestration against business purpose and accountability. The Trusted Living Twin™ is a way to make reviewed business knowledge and its relationships intelligible and reusable; it is not an agent, enforcement engine, or substitute for source systems. The proposed federated intake practice lets contributors and owners develop that knowledge in bounded increments, leaving an accepted client-owned record under agreed engagement terms. Broader live context delivery, independently operable client software, and automatic enforcement remain design or verification work, not results established by this paper.[13]

The purpose is to expand what the organization and its people can accomplish while maintaining trustworthy conduct. Leaders need to develop the AI capabilities, the human capabilities, and the relationship between them. That is the work of Non-Human Resources.

Source notes

1. Ethan Mollick, Management as AI superpower, January 27, 2026. Used for the discussion of direction, expertise, delegation, and evaluation. Co-intelligence is credited as an established contribution to the discussion, not claimed as a MortarCloud.ai invention.
2. Fabrizio Dell'Acqua and colleagues, The Cybernetic Teammate, Organization Science, published online June 12, 2026. A one-day P&G product-innovation workshop; the result concerns assessed proposals, not implemented business outcomes or long-term team replacement.
3. Jay R. Galbraith, Organization Design: An Information Processing View, Interfaces, 1974; Michael Y. Lee and Amy C. Edmondson, Self-managing Organizations: Exploring the Limits of Less-Hierarchical Organizing, Research in Organizational Behavior, 2017. Used at the scope of accessible primary abstracts and verified passages, not as evidence that the proposed structure has been tested.
4. Gartner, The AI Agents You Can't See Can Still Create Risk, September 30, 2026; Applying Uniform Governance Across AI Agents, May 26, 2026. Public analyst guidance on differing autonomy, access, and oversight arrangements.
5. Emma Wiles and colleagues, Putting AI on the Org Chart: Evidence on Delegation and Accountability, working paper, September 19, 2026. The stronger organization-chart subgroup was not the preregistered primary moderator; no result is presented here as proof that a particular governance mechanism resolves framing effects.
6. Amy Edmondson, Psychological Safety and Learning Behavior in Work Teams, Administrative Science Quarterly 44(2), 350-383, 1999; Ingrid M. Nembhard and Amy C. Edmondson, Making It Safe: The Effects of Leader Inclusiveness and Professional Status on Psychological Safety and Improvement Efforts in Health Care Teams, Journal of Organizational Behavior 27, 941-966, 2006; Jeeyoon Jeong, Byung-Jik Kim, and Julak Lee, The Effect of Job Insecurity on Knowledge Hiding Behavior: The Mediation of Psychological Safety and the Moderation of Servant Leadership, Frontiers in Public Health, March 13, 2023. These are bounded studies of team learning, leadership, and perceived insecurity, not universal causal findings about AI-driven workforce changes.
7. Judy Hanwen Shen and Alex Tamkin, How AI Impacts Skill Formation, arXiv version 2, February 1, 2026. Preprint; a short unfamiliar-library task and immediate assessment. Interaction patterns within the AI group were not themselves randomized.
8. Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, Generative AI at Work, Quarterly Journal of Economics, 2025. Final published version; a staggered rollout in one software company's customer-support operation.
9. NIST, Artificial Intelligence Risk Management Framework 1.0, January 26, 2023, GOVERN 2.1-2.3 and 3.2. Voluntary guidance; not a legal requirement or certification of this proposal.
10. Anthropic, Building Effective Agents, December 19, 2024, sections "What are agents?", "Agents", and "Combining and customizing these patterns". Practitioner guidance on architectural patterns, including supervised operation and composition; not a universal taxonomy or an assessment of MortarCloud.ai capabilities.
11. MortarCloud.ai, Master Plan v4.1, October 27, 2026, section 3; Capability and Evidence Register, October 30, 2026, sections 1-3; canonical framework structure and founder decision/provenance record supplied in the October 30 source handoff. Internal records read for this manuscript. Company capability descriptions reflect that documented state, rather than an independent audit of running software. Source locations and access scope are retained in the private evidence ledger.
12. Gartner, Scaling AI Means Closing the Risk Gap, October 5, 2026; The AI Value Gap Is Really an AI Risk Gap: Protecting Value at Scale, announced for November 4, 2026. The first is a published Gartner article, and the second is an upcoming webinar agenda as of November 2; no webinar conclusions or endorsements are claimed.
13. MortarCloud.ai, founder direction D06-D08 and D11, November 2, 2026, and proposed comprehensive business plan v4.2-R5, November 2, 2026. Private strategy source for Step Zero, bounded client-owned knowledge, federated contribution, flexible entry, and evidence-state restrictions. The founder's use of models interchangeably describes a working choice, not frictionless technical compatibility or identical capabilities. Approved strategic direction is not an independently tested deployment or acceptance of this manuscript.

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