ALIGN BY DESIGN

Scaling in the Age of AI

How Align by Design's Dilemmas Resolution Framework helps organizations redesign work roles and structure while they grow

Sirisha Bhamidipati & Harish Bhamidipati

Research report



Executive view

When organizations scale, the way work gets done has to change. A small team can depend on frequent conversations, flexible roles, and a few people making most of the decisions. As the organization grows, this becomes difficult to sustain. People need clarity on priorities, roles, decision authority, handoffs, measures, and accountability. Leaders also need to make sure that teams understand the larger objective and receive the context required to make sound decisions.

AI enters at a time when many organizations are still working through these scaling challenges. It changes the volume and speed of work. It also changes how information is created, how decisions are supported, and how work moves across functions. Assistants and agents can complete parts of a workflow that were earlier performed by people. This creates new questions about role boundaries, authority, control, and accountability.

Organizations therefore need to examine AI as part of organization design. Looking at a list of tasks and identifying what can be automated is insufficient. A role may include routine work, while also carrying customer understanding, judgment, relationships, expertise, accountability, and knowledge of exceptions. If leaders remove roles based on the automated portion of the work, they may lose contributions that were never formally documented.

The analysis should begin with the business outcome. Leaders need to understand what the organization is trying to achieve, what each workflow must deliver, and what contribution is expected from every role. They can then decide which activities an agent can perform, where AI can support a person, where human judgment is required, and who will remain accountable for the final outcome.

This will change the kind of people organizations need. Roles centred mainly on collecting information, coordinating routine activity, or processing standard work are likely to reduce. Domain expertise will become more important because someone must assess whether an AI generated answer makes sense in the specific business context. People in broader roles will need a credible area of depth. Learning agility will also matter because tools, workflows, and role expectations will continue to change.

Career development will need to change with this. A fixed career path may have limited value when roles are being redesigned frequently. Organizations need to make emerging roles and their capability requirements visible. Managers need to help people gain relevant experience and feedback. Individuals will carry greater responsibility for keeping their skills current and preparing themselves for the next relevant role.

The manager’s role will also shift. AI and digital workflows can take over parts of reporting, scheduling, tracking, routing, and routine review. Managers will need to spend more time helping people understand changing priorities, redesigning work, developing judgment, coaching employees, and resolving difficult exceptions. Reducing management layers without protecting these responsibilities can weaken the organization’s ability to develop people and handle complexity.

These changes create several dilemmas. Organizations must decide how much AI ownership should be centralized, how much autonomy teams and agents should have, where work should be standardized, how widely knowledge should be shared, and how much current productivity should be exchanged for future capability. Each choice offers benefits and creates costs.

Align by Design’s Dilemmas Resolution Framework helps leaders work through these choices. DRF recognizes that both sides of a dilemma may be valid. The appropriate choice depends on the business outcome, the context, the risks involved, and the organization’s ability to manage the consequences. AI adds more variables to these decisions and makes it even more important to choose deliberately.

Work should be redesigned before headcount decisions are made. Roles should be reviewed in terms of their full contribution to the organization. This gives leaders a stronger basis for deciding where AI should be used, how responsibilities should change, what skills will be needed, and how the organization should be structured for its next stage of growth.

Scaling changes the management system

The zero to one organization is held together by a small number of people who share context. Communication is direct. Priorities change through conversation. Work moves to whoever can solve the problem. Authority follows expertise, access to the founder, or urgency. Learning occurs through immersion because people see the whole problem and its consequences.

Those strengths become constraints during the one to ten journey. Information no longer reaches everyone at the same time. People interpret priorities through their function or location. Broad roles begin to overlap. Specialists improve the quality of individual activities and create new handoffs. Founders become decision bottlenecks. Managers emerge to translate priorities, allocate resources, integrate work, and supervise performance.

Scaling therefore changes the management system in six areas.

- Strategic context must move from the minds of a few leaders into language, priorities, measures, and repeated conversations that guide daily choices.
- Role clarity must define the outcome each role owns, the decisions it can make, the interfaces it depends on, and the issues it must escalate.
- Workflow must connect specialized contributions across functions and make handoffs, queues, exceptions, and dependencies visible.
- Control must move from direct observation to standards, measures, approvals, evidence, and review.
- Capability building must move from informal exposure to deliberate experiences, feedback, career moves, and professional communities.
- Leadership must move from direct involvement in every issue to creating the conditions in which other people can decide and perform.

Each response creates a dilemma. Strong standards improve reliability and can suppress useful local adaptation. Clear functional ownership builds expertise and can fragment customer outcomes. Delegation increases speed and can weaken consistency. Additional managers create coordination capacity and can lengthen the path of information. Career structures provide clarity and can preserve roles after the work has changed.

Greiner described recurring crises of leadership, autonomy, control, and coordination as organizations grow.[1] Later research explains why these tensions persist. Specialization improves defined tasks and increases the coordination needed across specialists. Hierarchy can coordinate interdependent work and reduce local autonomy and learning.[2] The goal of organization design is therefore to choose which tensions the business will carry and build mechanisms that keep those tensions productive.

AI adds a second scaling curve

AI changes the same management system that growth is already forcing the organization to formalize. It expands production capacity, lowers the cost of analysis, increases the volume of communication, and allows some decisions and actions to occur without continuous human involvement. It also makes the quality of data, context, instructions, and permissions part of day to day organization design.

The two scaling curves interact. A young company may be clarifying roles while AI unbundles the tasks inside those roles. It may be distributing authority while models prepare recommendations and agents execute actions. It may be creating functions to concentrate expertise while AI gives generalists access to plausible specialist output. It may be adding managers for coordination while technology absorbs information relay and creates a greater need for coaching and exception management.

Exhibit 1 AI changes the management system that organizations build for scale

Design area	What the one to ten journey already requires	What AI adds to the design problem
Communication and context	Proximity no longer keeps everyone aligned. Leaders must reinforce objectives, priorities, tradeoffs, and changes through timely channels and repeated conversations.	AI multiplies content and compresses response time. People and agents need a trusted context spine, clear provenance, and owners who refresh shared meaning as conditions change.
Role clarity and decision authority	Broad roles begin to overlap and founder decisions become queues. Each role needs clear outcomes, interfaces, decision rights, thresholds, and escalation.	Decision chains now include data, models, people, and agents. Higher speed and volume require policy, permissions, and exceptions because sequential approval cannot carry every case.
Data and knowledge	Systems of record and common definitions replace memory and local files as the organization grows.	Data becomes an input to recommendations and action. Quality, meaning, access, lineage, and refresh cycles become part of decision accountability.
Workflow and coordination	Specialization creates handoffs and dependencies. Processes, meetings, systems, and managers connect the work and resolve breakdowns.	Human and agent work must be orchestrated across prompts, models, data, tools, and teams. Routine coordination can be embedded while exceptions still need human ownership.
Roles and structure	Broad early roles become clearer and functions form around recurring capability. Career paths and reporting lines provide stability.	Tasks move between people and agents faster than job architecture changes. Role outcomes must stay clear while task bundles, capacity, and team boundaries evolve.
Expertise	Functions concentrate scarce knowledge, maintain standards, and create routes to mastery.	Expert quality output becomes widely available. Specialists shift toward standards, edge cases, validation, teaching, and the judgment required to detect plausible errors.
Control and accountability	Measures, approvals, reviews, and management layers replace direct observation as distance grows.	Controls move into workflow through permissions, thresholds, logs, and automated checks. Errors can travel quickly, so outcome ownership and stop rights must remain explicit.
Learning and careers	Repetition, supervision, and visible career paths help people build experience over time.	Skill requirements change faster and routine practice can disappear. Organizations must show next role requirements while individuals take greater ownership of relevance and learning.
Leadership and management	Managers coordinate work, relay information, motivate teams, monitor performance, and solve cross team problems.	AI absorbs parts of information relay and routine control. Managers spend more time preserving context, coaching, redesigning work, developing judgment, and handling exceptions.

The second scaling curve changes the speed, evidence, actors, and consequences inside familiar organization design choices.

The timing changes the risk of delay. In the earlier scale journey, an unclear role might create duplicated work or a slow escalation. With AI, the same ambiguity can be repeated through an automated workflow, multiplied across transactions, and hidden behind a fluent output. A weak decision path can become a fast decision path before the organization agrees who owns the result.

Data also changes status. In a traditional process, poor data creates a weak report. In an AI enabled process, data may shape a recommendation, a customer response, a prioritization decision, or an action taken by an agent. Data quality, semantics, access, provenance, and refresh cycles therefore sit inside decision authority. A

person who owns the decision needs enough visibility to understand the evidence and enough authority to change the workflow.

Evidence already shows the gap between individual use and organization redesign. A field experiment across 66 firms and 7,137 knowledge workers found that active users of generative AI spent about two fewer hours a week on email. The overall composition of work changed little.[3] Individual capacity improved while roles, workflows, and measures stayed largely intact.

The same gap appears in enterprise results. McKinsey's 2025 survey found that 62 percent of respondents were experimenting with AI agents, while 39 percent reported an enterprise level EBIT impact from AI.[16] BCG found that 4 percent of companies were creating substantial value and another 22 percent had created some value.[18] Operating model choices help explain the difference, including ownership, workflow redesign, talent deployment, data, and governance.[17]

AI can accelerate a weak design as easily as a strong one. It can produce more output for a workflow with unclear ownership. It can move information faster through an approval chain that has outlived its purpose. It can create local agents that solve a team problem and fragment the enterprise. It can remove entry work before the organization has created another route to judgment. The first question concerns the outcome and the system of work that produces it.

Design from outcomes before allocating work

An organization should decide the contribution of AI after it understands the contribution required from the role. This requires a clear chain from strategy to work.

Exhibit 2 The sequence for designing human and AI work

1	**Business outcome**	The customer or economic result the organization must achieve
2	**Performance outcomes**	What workflows teams and functions must deliver
3	**Role contributions**	The outcomes judgment relationships expertise decisions and execution each role must provide
4	**Task and decision allocation**	What people assistants and agents do at each step
5	**Design mechanisms**	Authority data standards evidence escalation control and learning
6	**Review triggers**	The changes in strategy performance risk capability or technology that reopen the design

Task automation is a downstream design choice. The outcome and role contribution come first.

A business outcome describes the customer, economic, risk, or mission result the organization needs. A performance outcome describes what a workflow, team, or function must achieve to support that result. Examples include reducing onboarding time, improving forecast accuracy, increasing first contact resolution, or shortening the time from insight to product change.

A role contribution describes the outcomes, judgment, decisions, relationships, expertise, and execution for which a role is accountable. Tasks sit inside that contribution. Some tasks can be automated. Some can be prepared by AI and reviewed by a person. Some need human involvement because the work depends on context, trust, consequence, negotiation, empathy, or learning.

This sequence exposes the difference between a productivity gain and an organization gain. A faster task may move the queue to the next team. A shorter analysis cycle may increase the volume of decisions waiting for approval. A good first draft may create more review work. The full workflow shows whether value improved and whether responsibility remained clear.

For each material workflow, leaders should define six things before changing the structure.

1. The business outcome and the measures that show whether it improved.

2. The performance outcomes required from each part of the workflow.
3. The role contributions needed to produce those outcomes.
4. The tasks and decisions that people, assistants, and agents will perform.
5. The authority, data, evidence, and escalation that support the allocation.
6. The learning experiences and review triggers that keep the design current.

The organization can then examine its dilemmas with a clear view of what the design must accomplish.

How the Dilemmas Resolution Framework applies

A dilemma exists when two legitimate needs pull the organization in different directions. Central authority protects enterprise coherence. Distributed authority places decisions near current information. Functional depth builds mastery and professional standards. End to end ownership reduces handoffs and clarifies responsibility for the result. Stable roles support accountability and identity. Fluid deployment moves skills toward changing priorities.

DRF asks leaders to choose a position that fits the strategy, work, maturity, talent, and risk. The choice should state what the organization expects to gain, what cost it accepts, and which mechanisms will prevent that cost from overwhelming the benefit.

AI changes a DRF discussion in five ways.

- It changes the capacity, speed, and economics on both sides of an existing dilemma.
- It introduces digital actors with goals, instructions, permissions, memory, and limits.
- It makes data and context direct inputs into decisions and execution.
- It changes how people gain experience, demonstrate proficiency, and move between roles.
- It connects local choices to enterprise exposure because a workflow can scale quickly.

A useful resolution records the outcome, the chosen position, the reasoning, the cost accepted, the supporting mechanisms, the owner, and the evidence that will trigger a review. This turns a call for balance into a design decision that can be tested.

The expanded dilemma landscape

The dilemmas of scaling remain. AI changes their intensity and creates additional tensions around data, human development, agent authority, and system control. The landscape below is a diagnostic for leadership teams. It helps them identify where choices are being made implicitly and where several choices may conflict.

Work and value

Enduring dilemma	Additional tension created or intensified by AI
Efficiency or effectiveness	Output can rise without improving the customer or economic result. Leaders must distinguish saved effort from value created across the full workflow.
Standardization or local adaptation	AI can personalize work inside a common process and can also spread unofficial local workflows at enterprise speed.
Automation or human involvement	The choice moves from whole jobs to specific decisions tasks exceptions and learning experiences inside a role.
Task optimization or end to end performance	A faster activity can increase queues reviews or risk elsewhere. Local productivity may weaken the complete outcome.

Enduring dilemma	Additional tension created or intensified by AI
Speed or quality	Higher output volume reduces the attention available for review and makes plausible errors more difficult to detect.
Scale efficiency or customer intimacy	Personalization can expand at low cost while consistency fairness brand voice and trust become harder to protect.
Current delivery or exploration	The cost of creating tests falls. Capacity can fragment across many experiments without a shared learning agenda.

Each pair contains two legitimate needs. The relevant position depends on the outcome and context.

Structure and authority

Enduring dilemma	Additional tension created or intensified by AI
Centralized or distributed decisions	Analysis and action can move closer to the work while enterprise exposure grows. Policy data ownership and case authority may sit in different places.
Functional depth or end to end ownership	Smaller outcome teams can cover a wider workflow and still need expert homes for standards difficult cases and development.
Tall or flat hierarchy	Information relay work declines while context translation coaching integration and exception ownership become more important.
Stable teams or fluid networks	People and agents can assemble around outcomes quickly while continuity belonging and accountable ownership become less obvious.
Stable roles or dynamic task allocation	Role outcomes need continuity even as tasks move frequently between people assistants and agents.
Central AI platform or domain ownership	Shared identity infrastructure and controls need enterprise ownership while workflow design and outcome accountability remain close to the domain.
Internal capability or external ecosystem	External models agents and partners add capacity while increasing dependency switching data confidentiality and knowledge risks.

Each pair contains two legitimate needs. The relevant position depends on the outcome and context.

People and capability

Enduring dilemma	Additional tension created or intensified by AI
Deep specialists or AI enabled breadth	Expert quality output becomes widely available while genuine mastery remains essential for validation edge cases and standards.
Hire or develop	New skills are scarce while current employees hold customer process and institutional context that technology needs.
Current productivity or future capability	Routine work disappears and can remove the practice through which future specialists managers and leaders build judgment.
Career paths or next role readiness	Static ladders age quickly. People need visibility into adjacent roles changing requirements and the evidence required to move.
Organization investment or individual ownership	Employers control access to work feedback and mobility while individuals carry greater responsibility for maintaining relevance.
Measured output or complete contribution	Visible production can rise while judgment teaching knowledge maintenance and exception handling disappear from performance measures.
Stable expertise or learning agility	Current mastery may depreciate faster. The organization must value the ability to learn release old routines and transfer judgment.

Each pair contains two legitimate needs. The relevant position depends on the outcome and context.

Leadership and management

Enduring dilemma	Additional tension created or intensified by AI
Coordination or coaching	Technology can absorb updates routing and routine follow up. Managers need to redirect capacity toward context development and judgment.
Enterprise direction or local interpretation	Strategic context must remain coherent while teams adapt it to changing customers workflows and agent behavior.
Control or trust	Workflow controls become stronger and more granular while excessive monitoring can weaken discretion ownership and responsible challenge.
Leader judgment or AI recommendation	Leaders gain faster analysis and can lose independent thinking when assumptions evidence and uncertainty remain hidden.
Current delivery or capability building	Pressure for immediate output can consume the coaching experiments and mobility needed to prepare people for changing roles.
Narrow spans or empowered wider spans	AI may support wider spans operationally while people development and complex exception work still constrain managerial capacity.

Each pair contains two legitimate needs. The relevant position depends on the outcome and context.

Knowledge and coordination

Enduring dilemma	Additional tension created or intensified by AI
Communication volume or shared meaning	AI can create summarize and personalize messages while priorities tradeoffs and common interpretation become harder to reinforce.
Persistent context or adaptive context	People and agents need stable reference points while strategy policy and operating conditions continue to change.
Proactive or reactive communication	Agents can initiate messages and actions before a person sees the need creating value and new risks of noise or misdirection.
Open knowledge or confidentiality	AI can retrieve combine infer retain and reproduce sensitive information across sources.
Codified knowledge or tacit judgment	More knowledge can be captured while context pattern recognition and experience remain essential for exceptions.
Structured or informal collaboration	Digital workflows improve visibility and can remove the conversations through which nuance trust and shared context travel.
Individual productivity or team coordination	Personal gains can increase handoffs reviews and downstream workload when the team workflow stays unchanged.

Each pair contains two legitimate needs. The relevant position depends on the outcome and context.

Governance and trust

Enduring dilemma	Additional tension created or intensified by AI
Experimentation or reliability	Low cost pilots can become critical operating systems before ownership evaluation monitoring and support are established.
Agent autonomy or human accountability	Systems can act at volume while people and the organization retain the consequence and may lack practical control.
Data driven or human judgment	Recommendations become easier to produce and harder to challenge when data meaning assumptions and uncertainty are opaque.
Transparency or performance	Complex systems may perform well on average while users struggle to explain or contest an individual outcome.

Enduring dilemma	Additional tension created or intensified by AI
Risk prevention or recovery	Some failures cannot be predicted. Faster action raises the value of containment reversibility fallback and stop rights.
Personalization or privacy and fairness	AI can tailor decisions at scale and create unequal treatment intrusive inference or inconsistent customer experience.
Central control or federated guardrails	Enterprise standards reduce exposure while domain teams need room to adapt controls to their workflow and consequence.

Each pair contains two legitimate needs. The relevant position depends on the outcome and context.

How to read the diagnostic

Start with a material business outcome and trace the complete workflow that produces it. The most consequential dilemmas usually appear where work repeatedly slows, fragments, escalates, loses quality, or depends on a few people who hold critical context. A dilemma also appears when two groups optimize legitimate goals that pull the full outcome in different directions.

Leadership teams can use the landscape in five moves.

- Mark the current position on each relevant dilemma. Describe how the organization actually works, including informal practices, local tools, and exceptions.
- State the benefit the current position creates and the cost it imposes. Evidence may include cycle time, rework, customer outcomes, control failures, employee mobility, or concentration of expertise.
- Apply the AI lens. Examine changes in capacity, decision volume, data dependence, speed, reach, learning, and consequence. Identify any agent that can recommend, decide, communicate, transact, or alter a record.
- Look for clusters. A choice about hierarchy may depend on choices about workflow visibility, decision authority, manager capability, and escalation. A choice about agent autonomy may depend on data quality, permissions, monitoring, and reversibility.
- Choose the few dilemmas that constrain the outcome most. Record the intended position, the costs accepted, the mechanisms required, the accountable owner, and the evidence that will reopen the choice.

The diagnostic becomes more valuable when different levels compare their answers. Senior leaders may believe authority is distributed while teams still wait for approval. A central function may believe knowledge is accessible while people rely on private files and personal relationships. These gaps reveal the lived organization design. They also show where adding AI could amplify ambiguity or where better technology could make a deliberate design workable.

Ten dilemmas deserve deeper attention because they shape several others. They affect the operating model, leadership work, capability building, and the boundary between human and digital contribution.

DILEMMA ONE

Standardize the work or preserve local adaptation

Why the tension appears at scale

Early teams adapt through direct conversation. People understand the customer, see unusual conditions, and change the approach in real time. As volume grows, variation becomes harder to distinguish. Some variation reflects customer need. Some reflects weak process discipline, missing knowledge, or repeated reinvention.

Standardization makes work repeatable. It supports quality, training, measurement, interoperability, and scale. Local adaptation protects judgment and responsiveness. Excessive standardization can remove the discretion that creates value. Excessive variation makes performance difficult to compare and useful learning difficult to spread.

What AI changes

AI can interpret unstructured information and tailor an output inside a standard workflow. This allows an organization to standardize the outcome, evidence, data, and controls while giving teams flexibility in the path. AI can also spread local variation quickly. A prompt, agent, or workflow built by one team may become an unofficial process used by hundreds of people.

The design question shifts from whether the entire process is standard to which elements require consistency. Customer facts, policy, quality thresholds, evidence, and escalation may need a common design. Language, sequence, or response may adapt within defined boundaries. The organization also needs a route through which useful local innovation becomes a maintained enterprise standard.

What a deliberate resolution requires

Classify variation by its source and consequence. Protect variation that improves the outcome. Remove variation caused by missing information or avoidable reinvention. Treat local AI workflows as operating assets once other people depend on them. Give a named owner responsibility for performance, change control, and retirement.

Additional questions leaders need to answer

- Which business outcome requires consistency across teams or locations?
- Where does local context materially change the right answer?
- Which elements must remain common for safety compliance data quality or interoperability?
- What may AI adapt and within which quality and authority boundaries?
- How will the organization detect unofficial prompts agents or workflows that have become critical?
- How will a useful local innovation become an enterprise standard?
- What level of variation signals fragmentation or control failure?

DILEMMA TWO

Centralize decisions or distribute authority

Why the tension appears at scale

In a small organization, decision authority often follows the founder or the person with the most context. Growth increases decision volume and spreads relevant information across functions and locations. The organization then needs explicit clarity on who owns the decision, who contributes, which thresholds apply, and what must be escalated.

Centralization creates coherence and allows a wider enterprise view. It can also create queues and distance from current customer or operating information. Distributed authority improves speed and local ownership. It can produce conflicting choices, repeated analysis, and unclear accountability.

What AI changes

AI makes analysis and recommendations available closer to the work. Agents can also execute an approved action. A single decision chain may now include source data, a retrieval system, a model, a business rule, a person, and an agent. Data quality and context quality become part of decision quality.

The speed and number of decisions also change. Layer by layer approval becomes impractical when thousands of recommendations or actions occur each day. The organization needs policy, thresholds, permissions, monitoring, and exception routes that allow distributed speed while preserving accountability. A leader remains accountable only when the leader can see the evidence, understand the boundary, and change the system.

What a deliberate resolution requires

Separate four forms of authority: setting policy, deciding an individual case, executing the action, and handling the exception. They can sit in different places. Central teams may own enterprise boundaries, data definitions, and risk tolerance. Local teams may own decisions within those boundaries. Agents may execute reversible actions and escalate cases that exceed confidence or consequence thresholds.

Additional questions leaders need to answer

- What outcome does the decision affect and who owns that outcome?
- Who sets policy and who decides an individual case?
- What data informs the decision and who owns its quality meaning and refresh cycle?
- Which decisions require enterprise context and which depend on local context?
- What may an agent recommend approve execute or reverse?
- Which thresholds send a case to a specialist or senior leader?
- Can the accountable person inspect and change the evidence logic permissions and escalation path?

DILEMMA THREE

Build functional depth or create end to end ownership

Why the tension appears at scale

Functions concentrate specialists, maintain methods, set standards, and create professional careers. They also divide a customer or business outcome across several teams. Each function can meet its own measure while the full workflow remains slow or inconsistent.

End to end ownership gives one team responsibility for the result and reduces handoffs. It can duplicate scarce expertise, weaken professional standards, and leave specialists without a strong development home. The tension grows as the organization adds products, locations, and customer segments.

What AI changes

AI gives smaller teams access to research, drafting, analysis, technical support, and institutional knowledge. Outcome teams can cover a wider part of the workflow. Access to fluent output can also create the appearance of mastery. Errors become difficult to see when the person reviewing the output lacks the underlying domain model.

Specialists therefore change their contribution. They may spend less time producing routine output and more time defining standards, curating knowledge, testing edge cases, reviewing high consequence work, and teaching others. Functions may become capability homes while outcome teams own priorities and the full result.

What a deliberate resolution requires

Assign one owner to the complete outcome. Place repeatable capability inside the outcome team where speed and context matter. Retain strong professional homes where expertise is scarce, errors are costly, or learning depends on a community of practice. Define how outcome leaders and capability leaders resolve conflicts over speed, standards, staffing, and development.

Additional questions leaders need to answer

- What complete customer or business outcome needs one accountable owner?
- Which handoffs currently delay dilute or obscure that outcome?
- Which capabilities can sit in the outcome team with AI support?
- Where can fluent output conceal weak expertise or missed exceptions?
- Which capabilities need a professional home because mastery risk or learning requires it?
- Who decides when outcome speed and professional standards conflict?
- How will specialists build mastery while working across outcome teams?

DILEMMA FOUR

Flatten the hierarchy or strengthen people leadership

Why the tension appears at scale

Hierarchy helps a growing organization translate priorities, allocate resources, integrate work, supervise performance, and escalate issues. It can also lengthen decision paths and create roles that mainly collect, summarize, and relay information. Wider spans reduce layers and may weaken coaching, development, and exception handling.

The relevant question concerns the work performed by each layer. Some management work exists because information is fragmented or measures are hard to see. Other work exists because people need context, feedback, development, conflict resolution, and judgment under uncertainty.

What AI changes

AI can handle parts of status collection, scheduling, reporting, documentation, routing, compliance checks, and routine review. Workflow systems can make progress and control visible without repeated management intervention. This reduces the value of layers built around information relay and administrative supervision.

The people leadership load can grow at the same time. Roles change faster. Employees need help interpreting expectations, building relevant capability, and learning from unfamiliar work. Human and agent contributions must be redesigned. Exceptions become a larger share of human work. Deloitte identifies coaching, work redesign, and the optimization of human and machine interaction as core future management capabilities.[14]

What a deliberate resolution requires

Assess each layer by the contribution it makes. Remove work that a clear workflow or shared evidence can perform. Preserve management capacity where teams face ambiguity, interdependence, development needs, or material consequences. Measure managers on team capability, decision quality, mobility, and outcomes alongside delivery.

Additional questions leaders need to answer

- What work does each management layer perform today?
- Which activities relay information or enforce a routine that technology can make visible?
- Which activities require coaching context integration conflict resolution or judgment?
- How many changing workflows roles stakeholders and agents can one manager understand well?
- If a layer is removed who develops people and owns difficult exceptions?
- What authority does a manager need to change an agent's goals data permissions or escalation?
- Which measures show whether a flatter structure improved outcomes while preserving capability?

DILEMMA FIVE

Keep stable roles or deploy fluid skills

Why the tension appears at scale

Stable roles provide accountability, pay logic, career identity, and a clear management home. They can also preserve task bundles after the work has changed. Fluid deployment moves people toward priorities and helps teams form around outcomes. Constant movement can weaken belonging, depth, continuity, and accountability.

During the one to ten journey, leaders often build career paths to show how people can progress through a growing structure. Those paths become less reliable when roles and tasks change faster than the job architecture.

What AI changes

AI changes the task mix within a role at different speeds. Two people with the same title may need very different capabilities because their workflows, customers, risks, or tools differ. The organization needs clarity on the enduring contribution of the role and flexibility in the skills used to deliver it.

Career development also shifts toward next role readiness. The World Economic Forum reports that employers expect 39 percent of workers' core skills to change by 2030.[10] Organizations need to make emerging roles and capability requirements visible. Individuals carry greater responsibility for maintaining relevance, seeking experiences, and demonstrating proficiency. Leaders still control access to meaningful work, feedback, mobility, and learning time.

What a deliberate resolution requires

Keep a stable home for accountability, pay, belonging, and development. Add a flexible contribution layer through projects, workflow assignments, rotations, and skill based staffing. Replace static career ladders with visible role adjacencies that show the skills, experiences, and evidence needed for the next relevant role.

Additional questions leaders need to answer

- Which accountabilities must remain stable as tasks change?
- Which tasks are shrinking growing or moving between people and agents?
- What is the enduring contribution of each role to the business outcome?
- Which next roles are becoming relevant and what evidence shows readiness?
- What learning responsibility sits with the individual manager capability home and enterprise?
- Who resolves conflict between a person's role home and temporary assignment?
- How will fluid deployment affect pay evaluation promotion and belonging?

DILEMMA SIX

Build deep expertise or expand AI enabled breadth

Why the tension appears at scale

Specialists develop domain models, pattern recognition, and professional judgment. Generalists connect disciplines and carry broader outcomes. A specialist model adds handoffs and can place scarce expertise in a queue. A generalist model can spread people beyond their competence.

Scaling organizations need both integration and depth. The choice depends on the consequence of error, the rate at which the domain changes, and the amount of judgment hidden inside apparently routine work.

What AI changes

AI gives people access to specialist quality language and analysis across many domains. A generalist can draft a contract clause, analyze data, prepare a forecast, or create a prototype. The output can be useful and still rest on a weak premise, stale context, or missed exception.

Research on the jagged frontier of AI found that performance gains depend on whether the task sits inside the system's capability boundary. Consultants performed better on tasks inside that boundary and were less likely to reach the correct answer on a task outside it.[4] Expertise increasingly includes knowing when an output deserves trust, what evidence is missing, and when escalation is required.

What a deliberate resolution requires

Build a workforce with deeper anchors and wider reach. Specialists should set standards, maintain knowledge, test difficult cases, and develop others. Broader roles should connect disciplines from a base of credible expertise and operate within clear authority limits. Tool use alone provides weak evidence of proficiency. Decisions and outcomes provide stronger evidence.

Additional questions leaders need to answer

- Which outcomes benefit from broader capability at the point of work?
- Where does shallow competence create material customer financial legal or safety risk?
- What expert judgment is required to evaluate an AI output?
- Which knowledge can be codified and which depends on experience?
- How will specialists spend time on standards difficult cases and teaching?
- What evidence will demonstrate proficiency beyond fluent production?
- What signals require escalation to deeper expertise?

DILEMMA SEVEN

Capture productivity or build future capability

Why the tension appears at scale

Scaling organizations need near term output and reliable economics. They also need a future supply of specialists, managers, and leaders. Training, supervision, and apprenticeship consume current capacity. Removing them can improve short term measures and weaken the future organization.

Routine work has historically served two purposes. It produces output and gives less experienced people repeated exposure to patterns, consequences, and feedback. The learning value is easy to miss when leaders assess the task only by current productivity.

What AI changes

AI often absorbs the first drafts, simple cases, research, and preparation through which people begin to learn a domain. A study of customer support agents found an average productivity improvement of about 14 percent, with larger gains for less experienced workers.[5] The same support can reduce independent practice when people accept the output without forming their own judgment.

The experience gap therefore widens. Deloitte found that 66 percent of surveyed managers and executives said most recent hires lacked full preparation, with experience the most common gap.[15] If organizations remove entry work and then hire only experienced people, the future supply of expertise becomes someone else's problem.

What a deliberate resolution requires

Allocate the productivity dividend explicitly. Use some released capacity for growth, service, or cost improvement. Use some for redesign and capability building. Replace lost practice with supervised exceptions, simulations, comparison of human and AI reasoning, rotations, and structured expert review. Measure proficiency growth alongside output.

Additional questions leaders need to answer

- Where will AI release capacity and how will the organization use it?
- Which current tasks build pattern recognition judgment or customer understanding?
- What learning disappears when those tasks are automated?
- What practice can produce equal or stronger capability?
- Which roles form the future supply of specialists managers and leaders?
- How much expert time is available for review coaching and feedback?
- Which measures track proficiency growth alongside cost cycle time and quality?

DILEMMA EIGHT

Move quickly with experiments or build reliable operations

Why the tension appears at scale

Experimentation helps a growing organization discover better ways of working. Reliable operations create repeatability, trust, and control. Heavy governance can stop useful tests. Unbounded experimentation can leave duplicate tools, fragile workflows, hidden costs, and obligations that no one owns.

The transition from experiment to operation is often unclear. A local tool may gradually become part of a critical workflow because colleagues find it useful. Dependence grows before ownership, support, and control catch up.

What AI changes

AI experiments are easy to start and difficult to see. Employees can create prompts, assistants, and agents without a formal project. The tool may use sensitive data, shape a decision, or communicate externally. Its behavior may also change when the model, data, prompt, or connected system changes.

The boundary between test and production therefore depends on reach and consequence. A private experiment with synthetic data carries a different exposure from an agent that updates records, sends messages, or influences a regulated decision. Reliability requires evaluation, monitoring, fallback, and an owner who can stop or change the system.

What a deliberate resolution requires

Create a graduated path from sandbox to live operation. Increase evidence, ownership, access control, monitoring, and approval as reach and consequence grow. Keep low risk learning easy. Treat production release as a business decision. Set review triggers for model changes, data drift, growing usage, and new dependencies.

Additional questions leaders need to answer

- What can teams test without prior approval and within which data boundaries?
- What change in reach consequence or dependency turns a test into an operating system?
- Who owns the business result when the tool enters live work?
- What evidence is required before the system reaches more users data or decisions?
- How will the organization discover informal AI use that has become critical?
- What changes require revalidation of the workflow?
- Who can stop retire or replace the tool?

DILEMMA NINE

Expand agent autonomy or preserve human accountability

Why the tension appears at scale

Delegation increases speed and capacity. Control protects the organization from actions that exceed authority or risk tolerance. Human delegation relies on roles, supervision, norms, and consequences. Agent delegation requires these mechanisms to be expressed through goals, permissions, data, thresholds, monitoring, and workflow.

Accountability weakens when a person is named as reviewer without enough time, expertise, evidence, or authority to intervene. The label remains human while practical control sits elsewhere.

What AI changes

An agent can interpret a goal, call tools, update systems, communicate with customers, and continue through several steps. Authority and accountability can separate. The agent acts while the organization carries the consequence. A review step offers limited protection when the volume exceeds human attention or the evidence is difficult to inspect.

Agent autonomy should therefore reflect consequence, reversibility, predictability, and observability. Low consequence work with clear feedback can support wider autonomy. High consequence actions need narrower permissions, meaningful human judgment, and a reliable stop mechanism. NIST emphasizes governance, measurement, monitoring, and risk management across the AI lifecycle.[7]

What a deliberate resolution requires

Define the agent's goal, action space, data access, spending or transaction limits, escalation rules, and stop rights. Name a human owner for the outcome and give that person the authority and evidence needed to own it. Monitor the full workflow, including whether people are accepting outputs without meaningful review.

Additional questions leaders need to answer

- What goal may the agent pursue and which actions may it take?
- Which actions are reversible and which create lasting customer financial legal or safety consequences?
- What information can the agent retrieve combine retain and change?
- Who monitors performance and who can stop the agent immediately?
- What evidence records what the agent saw decided changed and escalated?
- Does the accountable person have enough authority expertise time and system visibility?
- What level of volume makes human review ineffective?

DILEMMA TEN

Open access to knowledge or protect confidentiality

Why the tension appears at scale

Knowledge sharing reduces repeated work and helps teams make better decisions. Confidentiality protects customers, employees, intellectual property, negotiations, and regulated information. Tight restrictions can trap useful knowledge. Broad access can create exposure.

As the organization grows, knowledge also fragments across documents, systems, functions, and people. Search becomes harder while the need for shared context grows.

What AI changes

AI can make institutional knowledge available at the point of work. It can also combine information across sources, infer sensitive facts, retain material in a workflow, and reproduce it in a new context. The access question extends beyond who can open a document. It includes what a model or agent can retrieve, combine, generate, and send.

Knowledge quality matters alongside access. Stale policy, conflicting definitions, and undocumented exceptions can create confident answers that weaken shared meaning. Each important knowledge domain needs an owner, a refresh cycle, and a route for users to challenge the answer.

What a deliberate resolution requires

Classify knowledge by purpose and consequence. Give people and agents the minimum access required for the outcome. Separate retrieval from action. Record sensitive use. Assign owners to critical knowledge and build feedback into the workflow so context evolves when the business changes.

Additional questions leaders need to answer

- Which knowledge should be available across the organization and for what outcome?
- Which information requires restriction because of privacy regulation contracts or competitive value?
- Can the system combine individually permitted sources into a sensitive conclusion?
- Who owns the quality meaning and refresh cycle of critical knowledge?
- What data may be used for retrieval training evaluation and monitoring?
- Who can authorize an agent to send publish or act on information?
- How will useful learning be shared after confidential details are removed?

The dilemmas form one system

Each resolution changes the conditions of several others. Distributed authority increases the need for common context, data definitions, and evidence. Flatter structures increase the importance of manager quality and visible escalation. End to end teams increase the need for strong capability homes. Fluid deployment increases the need for stable accountability. Wider agent autonomy raises the value of identity, permissions, monitoring, and stop rights.

An organization can make a reasonable choice on every dilemma and still create an incoherent system. A company might flatten management, distribute AI ownership, widen data access, and increase agent autonomy at the same time. Each choice can improve speed. Together they can leave few people able to see enterprise exposure, maintain expertise, or develop inexperienced employees.

DRF therefore works at two levels. Leaders choose a position on each dilemma. They then test the combined design. Authority, information, capability, incentives, and accountability should meet where the outcome is produced.



Exhibit 3 A choice on one organization design dilemma changes the conditions of several others.

One coherent pattern for many scaling organizations contains five parts.

- Outcome teams own a customer or business result and the workflow that produces it.
- Capability homes maintain expertise, standards, careers, reusable knowledge, and learning experiences.
- A shared AI and data backbone provides identity, access, approved technology, monitoring, and common controls.
- Managers design the contribution of people and agents, preserve context, develop judgment, and handle exceptions.
- Governance sits inside the workflow through permissions, thresholds, evidence, escalation, and review triggers.

This is a reference pattern. Its form should follow the business model, risk, regulation, talent, and maturity of the technology.

A first principles process for leaders

The redesign should begin with one material outcome and the full workflow that produces it.

1. Define the business outcome. State the customer, economic, risk, or mission result and the limits the organization must respect.
2. Define performance outcomes. Clarify what each workflow, team, and function must achieve.
3. Map the work. Include decisions, tasks, handoffs, queues, exceptions, rework, information, and controls.
4. Define role contributions. State the outcomes, judgment, relationships, expertise, decisions, and execution each role must provide.
5. Identify the dilemmas. Locate the places where two legitimate needs compete and describe the benefit and cost on each side.
6. Apply the AI lens. Examine how assistants and agents change capacity, speed, decision volume, data dependence, risk, learning, and accountability.
7. Choose a position. Record the outcome sought, the benefits expected, the costs accepted, and the evidence behind the choice.
8. Design the mechanisms. Set decision rights, team boundaries, standards, permissions, measures, escalation, and learning practices.
9. Define the workforce implication. Determine the roles, skills, capacity, career moves, and leadership load required by the new design.
10. Set review triggers. Reopen the choice when performance, risk, capability, strategy, or technology changes.

Worker involvement improves implementation because employees can identify practical risks and changes in work that a central team may miss.[11] It also reveals informal coordination, judgment, and exception handling that rarely appear in a process map.

The measures should reflect the whole system. Cost and cycle time matter. Customer outcomes, quality, risk, exception rates, employee capacity, decision quality, and proficiency growth show whether value moved or the organization simply produced more activity.

Leadership when work is shared with AI

The earlier management model relied heavily on coordination, supervision, motivation, and control. Managers gathered information, translated direction, assigned work, checked progress, approved exceptions, and solved problems that crossed team boundaries. Much of that work emerged because information and workflow were difficult to see.

AI and digital workflows can absorb parts of information relay, scheduling, routing, progress tracking, documentation, and routine compliance. Leadership work becomes more concentrated in the areas that require human understanding and consequence.

Exhibit 4 Leadership shifts as operational coordination moves into the workflow

Leadership work	Earlier emphasis	Increased emphasis with AI
Context	Relay priorities and explain decisions from the top	Keep strategic context current and translate it into choices for people and agents
Work design	Assign roles resources and tasks inside an established process	Redesign outcomes roles and the allocation of work across people assistants and agents
Decisions	Approve cases and escalate issues through management levels	Set policy thresholds evidence and exception rights so decisions can move at greater speed and volume
Coordination	Run meetings collect updates and resolve handoffs	Remove dependencies orchestrate cross functional workflows and own the exceptions technology cannot resolve
Control	Review activity approve deviations and monitor measures	Embed controls in workflow inspect system signals and retain clear ownership for consequences
People	Guide motivate and evaluate performance in stable roles	Coach judgment build next role readiness create learning experiences and preserve trust through change

AI can absorb management activity. The leadership contribution grows in context judgment coaching system design and consequence.

The manager becomes the designer of a human and agent work system. The manager clarifies the outcome, explains changing context, sets the quality bar, decides where human judgment belongs, and ensures that exceptions reach the right expertise. The manager also protects the practice through which people develop judgment.

Coaching and mentoring therefore move closer to the center of the role. People need help understanding how their contribution is changing, which skills will keep them relevant, and what experiences will prepare them for the next role. Deloitte reports that effective development of direct reports can raise performance and increase the likelihood that people exceed their goals.[14] The value appears when managers have the time, capability, and measures to develop people.

Senior leaders carry a different set of responsibilities. They must keep the organization aligned around a changing strategic context, decide where digital autonomy is acceptable, make role and capability choices before setting headcount targets, and ensure that metrics and incentives reward work redesign. Microsoft found that organizational factors such as culture, manager support, and talent practices were more strongly associated with reported AI impact than individual factors in its 2026 survey of AI users.[13]

Leadership therefore operates on three time horizons.

- Deliver the current outcome through clear priorities, decision boundaries, and exception ownership.
- Redesign the work as technology changes capacity, risk, and the distribution of expertise.
- Build the people and institutional knowledge the organization will need next.

Removing coordination work should create leadership capacity. People still need context, coaching, and development across a wider span.

AI and the age of role redefinition

AI will automate a growing share of operational activity. That creates capacity and changes the economics of work. It also raises a deeper design question. What contribution should remain once routine production has reduced?

Every role should be examined as a contribution to a business outcome and to the performance of the workflow that produces it. Leaders can then separate five elements of the role.

- Work that an agent can execute within clear rules and permissions.
- Work that AI can prepare while a person applies judgment or accepts accountability.
- Work that depends on expertise, context, trust, negotiation, or empathy.
- Work that develops future judgment and needs deliberate human practice.
- Work that exists because the current process, structure, or system is poorly designed.

This analysis may reduce capacity in some areas. It may also change a function's mandate, move expertise closer to an outcome team, create a stronger capability home, or increase demand for scarce specialists. Each conclusion should follow the work analysis. A percentage reduction across functions can remove apprenticeship, exception handling, relationships, and institutional knowledge before leaders see their value.

The emerging workforce will have a different shape. Coordination only generalists will decline. Specialists who combine domain depth with AI fluency will become more important. The strongest broader profiles will connect disciplines from a base of credible expertise. Learning agility will determine how quickly people absorb new tools, release outdated routines, and build judgment for unfamiliar situations.

The responsibility for relevance is shared. The organization must make changing role expectations, next role opportunities, and proficiency standards visible. Managers must create access to experience and feedback. Individuals must take ownership of learning, seek assignments that build future value, and show evidence of capability. Career paths become a living map of possible moves and requirements.

Expectations of people rise in four areas.

- Frame the problem and define the outcome before asking AI to produce an answer.
- Exercise judgment when evidence is incomplete, consequences are material, or values conflict.
- Work across boundaries and integrate specialist perspectives into a coherent decision.
- Learn continuously through experiments, feedback, reflection, and exposure to exceptions.

DRF gives leaders a practical discipline for this redesign. It keeps both sides of a choice visible and connects technology choices to organization outcomes. It turns restructuring into informed decisions about outcomes, work, human contribution, capability, and the conditions under which people and AI perform together.

Sources

Research and institutional evidence

[1] Harvard Business Review. Evolution and Revolution as Organizations Grow. 1998. Open source

[2] Strategy Science. Hierarchies Knowledge and Power Inside Organizations. 2021. Open source

[3] National Bureau of Economic Research. Shifting Work Patterns with Generative AI. 2025. Open source

[4] Harvard Business School Digital Data and Design Institute. Navigating the Jagged Technological Frontier. 2023. Open source

[5] National Bureau of Economic Research. Measuring the Productivity Impact of Generative AI. 2023. Open source

[6] MIT Sloan School of Management. How AI Is Reshaping Workflows and Redefining Jobs. 2026. Open source

[7] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework Generative Artificial Intelligence Profile. 2024. Open source

[8] International Labour Organization. Generative AI and Jobs A Refined Global Index of Occupational Exposure. 2025. Open source

[9] Stanford Institute for Economic Policy Research. Canaries in the Coal Mine. 2025. Open source

[10] World Economic Forum. Future of Jobs Report 2025. 2025. Open source

[11] OECD. Using AI in the Workplace. 2024. Open source

[12] Anthropic. Economic Index New Building Blocks for AI Use. 2026. Open source

[13] Microsoft. Agents Human Agency and the Opportunity for Every Organization. 2026. Open source

Management and consulting research

[14] Deloitte. What Is the Future of Management. 2025. Open source

[15] Deloitte. Closing the Experience Gap. 2025. Open source

[16] McKinsey and Company. The State of AI in 2025. 2025. Open source

[17] McKinsey and Company. The Key to AI Value Is Hiding in Plain Sight Your Operating Model. 2026. Open source

[18] Boston Consulting Group. Where's the Value in AI. 2024. Open source