

What AI Changes About the First-Time Manager Role

When expertise is easier to access, leadership has to rest on something deeper.



1 transition

expert contributor to first-line manager

758 consultants

BCG AI capability-frontier experiment

5,179 support agents

workplace deployment across 3M+ chats

19 percentage points

lower correctness on an outside-frontier task

WHAT THIS REPORT EXAMINES

The first move into management has always required a shift from being valued for one's own output to being responsible for the conditions in which other people perform. AI changes that transition because some sources of expertise and coordination are becoming easier to distribute. Summaries, drafts, search, recommendations, and routine guidance can reach many people at once.

That does not make technical expertise less valuable. It changes what expertise can prove. A strong individual contributor may still know the work deeply, but access to a fast, fluent answer is no longer the same signal it once was. The new manager has to make value visible through judgment, standards, coaching, context, and accountable decisions.

The evidence also creates a second caution: assisted performance can improve faster than independent judgment develops. This report therefore treats readiness as a question of transfer. Can the person use AI without becoming dependent on it? Can judgment hold when the tool is uncertain, wrong, absent, or socially incomplete? Can the person stop rescuing the work and start developing the people who own it?

RESEARCH APPROACH | The report draws on AI task-performance experiments, workplace deployments, algorithmic-management research, and leadership-readiness theory. Direct studies that jointly track AI exposure, promotion into first-line management, and subsequent manager performance remain thin; where the evidence is indirect, BlackRidge treats the implication as evidence-informed rather than causal proof.

CONTENTS

The Decision Guide

00	Executive Summary	The first-time manager is losing an old advantage: being the person with the answers.
01	The Old Promotion Logic	Why expertise was a rational signal for management - and why it is becoming incomplete.
02	The Information Advantage Changes	How AI distributes knowledge, drafting, synthesis, and coordination across the team.
03	The Jagged Frontier	Why strong AI performance on one task tells you less than you think about the next one.
04	Borrowed Execution	How faster visible output can arrive before independent judgment develops.
05	The New Manager Role	What first-line managers must now frame, verify, coach, interpret, and own.
06	The Identity Transition	From answer person to accountable leader without abandoning expertise.
07	Develop the Right Evidence	How to build readiness through transfer, calibration, and real human consequence.
08	Organizational Responsibility	What must change in role framing, decision rights, support, and observation.
09	Research and Boundaries	What the evidence supports, what remains unproven, and what BlackRidge will not overclaim.

EXECUTIVE SUMMARY

The first-time manager is losing an old advantage: being the person with the answers.

For decades, technical expertise did double duty. It helped people perform well in the current role and gave newly promoted managers an informational advantage over the team. AI weakens that second advantage. It can distribute baseline knowledge, draft competent work, surface options, summarize performance data, and reduce some routine escalation. What remains is the harder part: deciding what matters, where the system may be wrong, how people will experience the decision, and who carries the consequence.



01

Expertise still matters

But the manager is less likely to be the only person with access to the answer. The value of expertise shifts toward setting standards and testing quality.

02

Performance can be borrowed

AI can accelerate novice output without proving that judgment has transferred. Fast work and mature leadership are different evidence.

03

Accountability does not automate

The manager still has to challenge, explain, correct, coach, and own consequential decisions when the tool cannot.

DECISION POINT

The question is no longer “Is this person the expert?” It is “Can this person govern work, judgment, and consequence when expertise is increasingly distributed?”

01

THE OLD PROMOTION LOGIC

We promoted expertise - then expected leadership.

That logic was not foolish. High performers understand the work, know the standards, solve difficult problems, and usually carry credibility with peers. For a first-time manager, those strengths can create an important foundation. The problem is that current-role expertise answers only part of the next-role question.

The old bargain

The expert performer was often the person who knew the process, remembered the exceptions, answered the difficult questions, and could rescue the work when others got stuck. Promotion converted that informational advantage into formal authority.

Do the work well

Know more than others

Become the answer source

Get promoted

What transfers

Credibility, craft, standards, pattern recognition, commitment to quality.

What becomes risky

Rescuing the work, protecting expert identity, equating answers with authority.

What must develop

Delegation, feedback, accountability, decisions under uncertainty, development of others.

Why AI matters

It weakens exclusive access to information - one of the old role's easiest sources of authority.

BLACKRIDGE FINDING

AI does not make technical excellence irrelevant. It makes technical excellence less sufficient as a proxy for leadership readiness.

THE INFORMATION ADVANTAGE CHANGES

AI distributes access to answers faster than organizations redesign the manager role.

The first-time manager historically had an easy way to establish value: know the work better, solve the problem faster, and answer questions the team could not. AI does not erase experience, but it changes who can produce competent output and how quickly.

Work element	Before widespread AI assistance	What AI can change	Managerial responsibility
Information	Knowledge accumulated through experience	Summaries, search, synthesis, retrieval	Verify meaning, missing context, and significance
Drafting	Expert creates first acceptable version	Competent first drafts become easier	Set standards, review quality, improve judgment
Coordination	Manager routes updates and follows up	Scheduling, prioritization, reminders, summaries	Resolve conflicts, trade-offs, and exceptions
Coaching	Manager supplies answers and examples	AI can provide baseline guidance	Develop thinking, calibrate judgment, preserve trust
Decision support	Manager integrates available evidence	AI can generate options and recommendations	Own the choice, explain it, and absorb consequence

The shift is subtle but consequential: the manager's value moves away from being the team's information router and toward becoming the person who decides what information can be trusted, how it fits the context, and what should happen next.



ROLE SHIFT

The strongest individual contributor may enter management with deep expertise - but not with the same exclusive informational advantage. That changes the source of authority, not the need for leadership.

THE JAGGED FRONTIER

AI can raise performance and still reduce judgment quality.

There is no general rule that “AI makes managers better.” In a preregistered study of 758 BCG consultants, AI users completed more tasks, worked faster, and produced higher-quality work on tasks inside the model’s capability frontier. On a complex managerial problem outside that frontier, AI users were 19 percentage points less likely to reach the correct answer. [3]



For a first-time manager, the implication is behavioral. Confidence has to become calibrated: strong enough to act, but responsive to task fit, stakes, uncertainty, and reversibility. The manager must know when to use AI, when to verify it, and when to slow the decision down.

MANAGERIAL LESSON

Fluency is not validity. A ready manager is not the person who can get the best AI output fastest; it is the person who can recognize when the output should not be trusted.

LEARNING VERSUS PERFORMANCE

Borrowed execution can look like capability.



A real customer-support deployment followed 5,179 agents across more than three million chats. AI recommendations increased productivity, and less-experienced workers benefited most. Tenure curves suggest AI-assisted agents reached a performance benchmark after roughly two months that unaided agents reached after eight to ten months. [2]

2 months

8-10 months



What the study could not tell us

The researchers could not determine whether lower-skill workers had learned the underlying reasoning independently or were following AI recommendations successfully. That difference becomes critical when the next case is novel, ambiguous, or socially consequential.

A Execution compression

AI can make visible output arrive earlier.

B Judgment development

Readiness requires transfer when the tool is absent or wrong.

BOUNDARY

Faster execution is evidence about performance on a task. It is not evidence that a new manager matures faster in conflict, fairness, accountability, or judgment.

02

THE NEW MANAGER ROLE

From answer source to accountable leader.

If AI can supply information, draft a response, recommend an action, and summarize performance, then the manager's value has to become visible somewhere else. The first-line role increasingly depends on six responsibilities that cannot be inferred from technical output alone.

WHY THESE RESPONSIBILITIES MATTER

These are not a new list of AI competencies. They are core management responsibilities that become easier to see once information and routine coordination are less scarce. The manager still has to define the work, decide what evidence deserves trust, set the standard, and translate a recommendation into action that people can understand and accept.

That is why tool fluency is weak evidence of readiness by itself. A candidate can be highly capable with AI and still avoid conflict, over-trust polished output, fail to develop another person's judgment, or hesitate to own a consequential choice. The role shifts from possessing answers toward stewarding decisions and conditions.

Frame the work

Define what problem is actually being solved and what standard matters.

Calibrate trust

Know when AI is useful, uncertain, incomplete, or outside its frontier.

Set standards

Judge whether output is merely fluent or actually fit for purpose.

Coach thinking

Develop another person's judgment instead of becoming the permanent answer key.

Read the room

Notice fear, silence, unfairness, resistance, and context absent from a dashboard.

Own outcomes

Explain, correct, escalate, and personally carry consequential decisions.

ROLE FINDING

The first-time manager no longer earns authority mainly by possessing answers. Authority increasingly comes from judgment, clarity, fairness, follow-through, and the willingness to own consequence.

WHAT SHOULD BECOME VISIBLE

Six readiness signals for an AI-enabled first-line role

These are not a stand-alone AI competency scale. They are existing leadership-readiness behaviors expressed in a changed operating environment. The evidence is strongest for ownership, calibration, pattern detection, resource judgment, and the continued importance of interpersonal interpretation.

WHY THESE SIX SIGNALS

AI-assisted work can hide errors and dependency behind fluent output. That makes six existing readiness behaviors especially revealing: ownership, calibrated confidence, pattern awareness, self-monitoring, interpersonal reading, and resource judgment.

These are not new AI competencies. They are an evidence-informed emphasis within BlackRidge's broader readiness model. Direct validation in AI-enabled first-line roles remains limited, so they are priorities to observe and test.


Readiness signal	What it looks like in an AI-enabled role	What to observe
Owning the Outcome	Traces evidence, explains the decision, corrects errors, carries consequence.	Can the person defend or reverse an AI-assisted people decision?
Decision Confidence	Acts without worshipping certainty; changes reliance with stakes and task fit.	When do they check, escalate, delay, or override?
Pattern Awareness	Looks for repeated errors, group disparities, missing context, distribution shifts.	Can they detect a pattern beyond one wrong output?
Inner Awareness	Notices deference, reflexive distrust, confirmation seeking, cognitive offloading.	Can they explain why they trusted or rejected the recommendation?
Reading the Room	Sees fear, silence, surveillance concerns, fairness, legitimacy.	What human signal is missing from the dashboard?
Resource Confidence	Routes the problem among AI, data, peers, HR/legal, escalation, conversation.	Can they justify who or what should carry each part?

READINESS SHIFT

The relevant question is not "Can this person use AI?" It is "Can this person remain accountable, socially aware, and appropriately skeptical while AI is part of the work?"

THE PEOPLE DECISION

The AI says the employee is underperforming. Now what?



Imagine a first-time manager receives an AI-generated performance summary showing one employee below target. The system highlights missed deadlines, lower output, and a pattern of delayed responses. A weak manager either accepts the dashboard as neutral fact or rejects it because “AI cannot understand people.” A ready manager works the decision.

01

Reconstruct the evidence

What inputs drove the summary? What was measured?

02

Look for missing context

Was work reassigned? Was the employee covering a gap?

03

Test the pattern

Is this isolated, repeated, or different across comparable employees?

04

Hear the person

What context, explanation, or fairness concern is invisible to the system?

05

Own the decision

Can the manager explain the final call without citing the tool as authority?

The point is not to force every people decision through a five-step ritual. It is to make the governing behaviors observable: traceability, context, pattern recognition, employee voice, and accountable ownership.

DECISION POINT

A first-time manager should never be able to say, “The system made the decision.” AI can inform the judgment; the leader must still make the judgment.

FROM EXPERT TO LEADER

The hardest shift may be identity, not technology.

The expert identity is powerful because it is earned. It says: I know the work. People come to me. I can solve the problem. AI can destabilize that identity by making competent answers available to more people more quickly. The first-time manager then faces a choice: compete with the tool and keep proving expertise, or reorganize expertise around leadership.

RETAIN

Credibility and craft

Keep the standards, experience, judgment, and pattern recognition that make leadership believable.

RELEASE

Answer-source identity

Loosen the need to be the smartest person in the room, rescue every task, or prove value by supplying the answer.

BUILD

Decision stewardship

Use expertise to frame the work, test outputs, coach thinking, and carry the consequence through others.

The identity trap

When the manager feels uncertain, the old role offers immediate relief: reopen the file, rewrite the draft, answer the technical question, fix the work personally. AI can make that pull stronger because it keeps the manager close to production. Readiness means using expertise without letting expertise become an escape route from leadership.

IDENTITY FINDING

The first transition is not about abandoning expertise. It is about refusing to use expertise - or AI fluency - as a substitute for authority, development, and accountability.

DEVELOPMENT DESIGN

Develop transfer, not just tool fluency.

If AI changes where answers come from, leadership development should not stop at prompting, productivity, or tool adoption. The stronger design makes judgment visible under changing levels of AI support.

WHAT TRANSFER MEANS

Tool fluency shows that someone can operate a resource. Transfer asks whether the underlying reasoning survives a change in conditions. If AI is removed, gives a weaker answer, or produces a recommendation that conflicts with local context, can the person still define the problem, identify the standard, compare alternatives, and act?

Leadership development should therefore vary the level of AI support rather than maximize it. Practice should include situations where AI is useful, situations where its capability is uncertain, and situations where human context changes the decision. The aim is not independence from AI; it is independence of judgment.

Development experience	What it should reveal	Why it matters
Paired AI / no-AI cases	Can the person perform, explain, and adapt without permanent tool dependence?	Separates assisted output from transferable judgment.
Inside / outside frontier cases	Does reliance change when the task becomes less suitable for AI?	Tests calibrated trust rather than generic confidence.
Systemic pattern exercise	Can the person detect repeated bias or missing context across several outputs?	Moves beyond spotting one obvious error.
Employee voice scenario	Can the person hear fairness and legitimacy concerns that are not in the data?	Preserves interpersonal reading under automation.
Resource-routing decision	Can the person choose among AI, data, peers, HR/legal, escalation, and direct conversation?	Tests practical resource confidence.
After-action reflection	What changed after an AI-assisted failure or surprise?	Distinguishes learning behavior from tool enthusiasm.

Development principle: do not make AI the curriculum. Make AI one of the conditions under which leadership judgment must hold.

DEVELOPMENT FINDING

The goal is calibrated resource use - neither overreliance nor reflexive rejection. A ready manager should be able to use AI well and still think without it.

MAKE READINESS OBSERVABLE

Do not wait for the title to discover whether the role changed.

Organizations can create bounded experiences before promotion that expose the new role directly. The objective is not to stage a hidden test. It is to let candidates practice consequential leadership, receive feedback, and show whether judgment transfers when AI is present, absent, incomplete, or persuasive-but-wrong.

Applied experience	What it can reveal	What to observe
Review an AI-generated status summary	Verification, pattern detection, escalation judgment	What does the candidate check before acting?
Coach an employee using AI-assisted work	Developmental attention, standards, ownership transfer	Does the candidate coach judgment or simply correct output?
Make a bounded people decision	Fairness, evidence calibration, accountability	Can the rationale survive challenge without citing the tool?
Lead a cross-functional AI-use disagreement	Peer navigation, decision rights, resource judgment	Can the candidate create clarity without formal authority?
Run a no-AI transfer case	Independent reasoning and learning transfer	Does the person retain the judgment when assistance disappears?

Evidence should be behavioral

Ask what the candidate checked, challenged, delegated, explained, corrected, or followed through - not whether the person "used AI effectively."

Evidence should remain developmental

Candidates should know what responsibility is being practiced, receive feedback, and get another chance to improve before a conclusion is drawn.

EVIDENCE FINDING

Readiness should be demonstrated in work that resembles the role ahead - including the conditions created by AI - rather than inferred from reputation built in different work.

THE ROLE MUST BE DESIGNED

A ready candidate can still enter an unclear AI-enabled role.

Leadership readiness is not solely the individual's burden. Organizations have to define where AI is allowed to assist, which decisions remain human, how appeals work, what evidence is retained, and what authority the new manager actually holds.

Organizational obligation	What good design looks like	What failure creates
Define AI decision rights	Specify where tools may draft, recommend, rank, monitor, or decide.	The new manager inherits ambiguous responsibility.
Name non-delegable decisions	Clarify people decisions that require explicit accountable human judgment.	"Human review" becomes a ceremonial approval click.
Preserve employee voice	Create channels for context, appeal, correction, and explanation.	Efficient decisions lose legitimacy and trust.
Protect learning opportunities	Give emerging managers low-stakes chances to decide, coach, and recover.	AI handles the routine work while judgment remains underdeveloped.
Observe real behavior	Review conversations, reasoning, verification, and follow-through.	Feedback focuses on output quality instead of leadership process.
Support the first 90 days	Align manager, sponsor, HR, role expectations, and AI-use boundaries.	The new leader has responsibility without a workable operating system.

Questions the organization should answer before Day 1

What decisions will the manager own? Where may AI recommend but not decide? What evidence must be retained? Who can appeal? What human context must be heard? Who coaches the first difficult people decision? How will readiness be reviewed at 30, 60, and 90 days?

ORGANIZATIONAL FINDING

Do not ask a first-time manager to prove readiness inside an AI-enabled role whose decision rights, accountability, and support have never been made explicit.

RESEARCH AND BOUNDARIES

What this report is - and is not.

This report translates current evidence on generative AI, algorithmic management, task performance, critical thinking, and human oversight into a focused question about the first leadership transition. Direct studies that jointly measure AI exposure, readiness for promotion, and later first-line-manager performance remain largely absent.

It is	It is not	Strongest evidence	Important unknown
A research-informed interpretation of how AI changes the context in which first-line readiness is expressed.	A validated AI leadership-readiness scale or proof that AI causes better or worse managers.	AI assists coordination and can improve or harm performance depending on task and reliance.	Whether AI changes promotion timing, spans of control, or new-manager time-to-proficiency causally.

Publication guardrails

- Do not treat software-mediated work as full automation unless a study demonstrates it.
- Do not equate faster output with mature judgment or readiness to lead.
- Do not claim AI is eliminating first-line managers; direct causal evidence is absent.
- Label consultancy and vendor evidence as directional rather than causal.
- Treat BlackRidge readiness implications as evidence-informed where direct validation is absent.

Selected research

- [1] OECD (2025). Algorithmic Management in the Workplace.
- [2] Brynjolfsson, Li & Raymond (2025). Generative AI at Work. QJE / NBER 31161.
- [3] Dell'Acqua et al. (2023/24). Navigating the Jagged Technological Frontier.
- [4] Cecil et al. (2024). Incorrect AI advice and explainability. Scientific Reports.
- [5] Köchling et al. (2024). Algorithmic bias in hiring. Frontiers in Psychology.
- [6] Lee et al. (2025). The Impact of Generative AI on Critical Thinking. CHI.
- [7] Hoffmann et al. (2024). Generative AI and the Nature of Work. HBS WP 25-021.
- [8] Capgemini Research Institute (2024). Gen AI at Work.

USE BOUNDARY

The strongest defensible conclusion is not that AI replaces the first-time manager. It is that AI changes the work through which readiness must now be demonstrated.



THE FIRST LEADERSHIP TRANSITION

The first-time manager no longer has to be the person with the best answer.

They do have to be the person who can decide what the answer means - and own what happens next.

AI changes where information comes from and how quickly competent output can be produced. It does not absorb the responsibility to set standards, develop people, resolve conflict, interpret context, make consequential decisions, and carry the outcome.

BUILD READINESS FOR THE ROLE AHEAD - NOT THE INFORMATION ADVANTAGE BEHIND.

blackridgeinstitute.com/first-time-managers/

blackridgeinstitute.com/insights/ai-leadership-readiness/