



REWIRING LEARNING & DEVELOPMENT FOR THE AI ERA

What enterprise AI strategy tells us about the future of organizational capability

The central question for L&D is no longer how to create learning faster. It is how to help the organization build capability fast enough to keep up with the work.

DR. RAVINDER TULSIANI

Enterprise Capability Executive

rtulsiani.com



ABOUT THIS REPORT

A strategy paper for L&D leaders facing a different kind of capability problem

This report translates current enterprise AI strategy and research from MIT, Stanford and Google into an L&D operating model. It does not treat AI as a content-production tool or as a topic for another curriculum. It examines what changes when AI alters tasks, workflows, expertise, performance support and the location of capability itself.

Design intent

The report uses a consulting-style editorial system: high signal density, restrained color, generous white space, explicit exhibits and short decision-oriented sections. The visual identity is aligned to the Ravinder Tulsiani executive brand: capability systems, AI readiness and measurable learning impact.

Evidence base

The source foundation includes nine enterprise AI strategy resources plus supporting Stanford research on generative AI at work and human-AI evaluation. Full source notes appear at the end.

Important limitation

MIT Sloan Management Review's GenAI Strategy Toolkit is a paid proprietary product. Only its publicly available description and framework components are reflected here.

CONTENTS

The argument in one line

AI changes not only what people need to learn, but where capability should reside and how organizations should build it.

Executive summary	04
01 AI exposes a problem L&D already had	06
02 The task becomes a more useful unit than the job	08
03 Stop assuming all capability has to live inside the employee	10
04 Expertise becomes more valuable, not less	11
05 Learning needs to move closer to work	13
06 L&D needs a different response to speed	15
07 The minimum effective intervention	16
08 Replace the learning portfolio with a capability portfolio	17
09 Measure the human-AI system, not the technology in isolation	19
10 AI literacy is necessary, but broad AI training is not a strategy	20
11 Governance should appear where decisions are made	21
12 The instructional designer becomes a capability designer	22
13 L&D needs its own AI operating model	24
14 A practical maturity model	27
15 What L&D should do now	28
16 The leadership question	29
Conclusion	30
Source foundation	31

EXECUTIVE SUMMARY

AI is forcing L&D to answer a larger question

Most Learning & Development teams are approaching AI from the wrong end of the problem.

They are asking how AI can help them create courses faster, personalize learning, generate assessments, build simulations or improve instructional design. Those are useful questions. They are also too small.

AI is beginning to change the work that L&D exists to support. Tasks are being automated, expertise is becoming easier to access, employees are working alongside increasingly capable systems, and the shelf life of some skills is shrinking. In that environment, improving the efficiency of the existing L&D model will not be enough.

The more important question is: How does an organization build and adapt capability when both the work and the technology supporting that work are changing continuously?

Recent work from MIT, Stanford and Google offers a surprisingly consistent answer. Successful AI adoption depends less on access to powerful technology than on the organization's ability to redesign work around it. Companies need to change workflows, clarify where people add value, build new skills, capture expertise, experiment with new ways of working and measure whether the combination of people and AI performs better.

Stanford's study of 51 enterprise AI implementations is particularly revealing. Organizations using similar technologies for similar purposes produced dramatically different outcomes. According to the researchers, the difference was not the underlying AI model. It was organizational readiness, processes, leadership and willingness to change.

That finding should matter enormously to L&D. It means the constraint on AI transformation may increasingly be organizational capability rather than technology.

But L&D only inherits that opportunity if it expands its definition of the work.

The future L&D function cannot primarily operate as a producer and distributor of learning. It needs to become part of the system through which the organization identifies capability gaps, redesigns work, develops people, embeds support, captures expertise and learns whether interventions are working.

The strategic shift

From courses to capability. From jobs to tasks and workflows. From knowledge transfer to performance enablement. From curriculum cycles to continuous adaptation. From completion metrics to performance evidence.

EXECUTIVE SUMMARY

Six shifts define the emerging L&D operating model

EXHIBIT Six shifts for L&D in the AI era

Shift	What changes
Courses → capability	Start with the business outcome and the required performance, not the requested learning asset.
Jobs → tasks and workflows	Analyze where AI changes the work at the task level and what the human contribution should become.
Knowledge transfer → performance enablement	Decide what must live in the person, the workflow, the technology and the team.
Broad programs → high-value use cases	Concentrate resources on a small portfolio of capability problems with material business consequences.
Completion metrics → performance evidence	Measure the person + technology + workflow system, not learning activity in isolation.
Content production → capability design	Move instructional design capacity toward diagnosis, judgment, experimentation and evidence.

The result is not the end of training. It is a clearer understanding of when training is actually the right answer.

AI exposes a problem L&D already had

AI did not create L&D's fixation on content. It exposed it.

When an organization identifies a performance problem, the request that reaches L&D often arrives already packaged as a solution: "We need training." "We need an e-learning module." "We need everyone to understand this." "We need a course on AI."

Once the request is framed that way, much of the real diagnosis has already been skipped.

MIT's guidance on AI strategy starts somewhere else. Organizations are advised to begin with critical business problems, determine where AI might create value, examine the data and capabilities required, and then build the necessary organizational response.

L&D should apply the same discipline to performance.

1	What do people need to do?
2	Why is that not happening now?
3	What is getting in the way?
4	What has already been tried, and what happened?
5	What evidence would tell us the problem has improved?

01 AI EXPOSES A PROBLEM L&D ALREADY HAD

These questions sound basic because they are. They are also routinely bypassed.

A manager who avoids difficult feedback may know exactly how to conduct the conversation. Another course on feedback will not solve fear of conflict, weak accountability or a culture where candour carries personal risk.

An employee making frequent processing errors may not lack knowledge. The process may be unnecessarily complex.

A salesperson struggling to prepare for meetings may not need product training. They may need better access to customer information.

A new employee overwhelmed by policy may not need to memorize more policy. They may need reliable guidance at the moment a decision is made.

AI makes this distinction harder to ignore because it introduces additional intervention choices.

EXHIBIT The intervention set is now wider than training

If the constraint is...	Possible response
Knowledge or skill	Learning, practice, coaching
Information access	Search, retrieval, performance support, AI assistance
Process complexity	Workflow redesign, simplification, automation
Weak reinforcement	Manager action, feedback, incentives
Decision quality	AI augmentation, expert review, structured judgment

Training should be one possible response to a capability problem, not the default response.

The task becomes a more useful unit than the job

Much of workforce development is still organized around roles. We create competency models for managers, learning pathways for salespeople, academies for leaders and curricula for functions.

AI complicates that structure because it rarely changes an entire job at once. It changes tasks within the job.

Stanford economist Erik Brynjolfsson argues that organizations should break roles into their "atomic unit of work," the task, then determine which tasks should remain human-led, which can be automated and which are best performed by people working with AI.

That idea has direct consequences for L&D.

Consider an instructional designer. The role may include interviewing stakeholders, analyzing performance problems, researching a subject, structuring content, writing learning objectives, drafting scenarios, producing graphics, developing assessments, reviewing SME feedback, testing learning and analyzing outcomes.

Treating "instructional designer" as one capability obscures what is actually happening.

02 THE TASK BECOMES A MORE USEFUL UNIT THAN THE JOB

Some tasks can increasingly be delegated to AI. Some can be accelerated by AI. Some become more important because AI can perform the surrounding work.

Stakeholder diagnosis, judgment, context, experimentation and evaluation may increase in value while first-draft production decreases.

The same pattern applies elsewhere. A financial analyst may spend less time assembling information and more time interpreting exceptions. A manager may spend less time drafting documentation and more time making difficult judgment calls. A customer service representative may retrieve fewer answers manually while taking greater responsibility for nuance, escalation and empathy.

This changes workforce development. Instead of asking, "What capabilities does this role need?" L&D should increasingly ask: "What work actually happens inside this role, how is AI changing that work, and what should the human contribution become?"

EXHIBIT A practical task classification

Category	Decision rule
Human-led	Keep the task primarily human when success depends heavily on judgment, accountability, trust, empathy, contextual interpretation or values.
AI-assisted	Use AI to improve the quality, speed or range of human performance while keeping a person responsible for the outcome.
AI-led with human oversight	Allow AI to perform most of the task while a person reviews exceptions, verifies quality or retains accountability.
Automated	Remove routine work where human involvement adds little value and the risk is understood.

Implication for L&D

Every task category creates a different development requirement. Job-level competency models are no longer enough.

Stop assuming all capability has to live inside the employee

For decades, L&D has operated under an implicit assumption: if people need to perform something, we should teach them enough that they can carry the capability with them.

That assumption made sense when access to information was slow and performance support was limited. It makes less sense when an employee can access relevant information, AI guidance, examples, coaching and analysis inside the workflow.

The design question changes. For any performance requirement, determine where the capability should reside.

EXHIBIT Where should capability reside?

Location	Best suited for
In the person	Foundational knowledge, fluency, judgment, safety-critical capability and skills that must be available without assistance.
In the workflow	Infrequent decisions, procedural guidance, checklists, prompts and contextual support.
In AI	Retrieval, summarization, drafting, pattern detection, option generation and first-pass analysis.
In the team	Social expertise, coaching, shared judgment, escalation and specialist knowledge.

Develop only as much individual capability as the work actually requires. Put the rest where people can reliably access it.

Expertise becomes more valuable, not less

One tempting conclusion from generative AI is that expertise matters less because knowledge has become easier to obtain. The evidence points in a more complicated direction.

A Stanford field study involving 5,172 customer support agents found that employees given access to generative AI assistance increased productivity by 15% on average. The strongest gains went to less experienced and lower-skilled employees. Experienced high performers gained much less.

Why? The AI system was effectively making some of the patterns embedded in expert performance available to less experienced employees.

That creates a major opportunity for L&D. Organizations contain large amounts of expertise that rarely appear in formal training: what experienced employees notice first, the questions they ask, exceptions they recognize, trade-offs they make, signals they distrust, when they escalate, and why they reject an apparently reasonable answer.

Traditional knowledge management has struggled to capture this. AI makes part of that problem more tractable.

But there is an important catch. If organizations use AI to reduce their reliance on experts without continuing to develop expertise, they risk consuming the very resource their systems depend upon.

04 EXPERTISE BECOMES MORE VALUABLE, NOT LESS

Brynjolfsson makes a related argument when discussing the future of apprenticeship. Senior experts increasingly become sources of high-quality judgments, standards and "gold answers" that improve the performance of both junior workers and AI systems.

L&D therefore needs a more deliberate expertise strategy.

The question is no longer just: "How do we train beginners?" It is also: "How do we capture expert judgment without destroying the conditions that produce experts?"

EXHIBIT The expert's role shifts

Less emphasis on	More emphasis on
Repeating routine answers	Reviewing difficult cases
Holding undocumented knowledge	Explaining reasoning and exceptions
Doing all high-quality work personally	Defining quality standards
Serving as the only escalation path	Coaching and system improvement
Protecting expertise from automation	Improving the knowledge and AI system

Junior employees can reach useful performance sooner. Experts move toward increasingly difficult judgment. The organization's knowledge system improves as both groups work.

Learning needs to move closer to work

The standard learning model separates development from performance. A person attends a program. They understand the material. They return to work. Then we hope understanding survives contact with reality.

Sometimes it does. Often it does not. Pressure, ambiguity, competing priorities, weak manager support and environmental constraints interfere.

AI creates an opportunity to shorten the distance between learning and application.

EXHIBIT Two learning logics

Traditional separation	Embedded performance
Learn	Attempt
Remember	Receive support
Apply	Act
Hope for transfer	Receive feedback
	Improve

A manager preparing for a difficult conversation might use an AI coach to rehearse objections immediately before the meeting. A salesperson could receive account-specific guidance while preparing for a customer discussion. A technician could access troubleshooting assistance at the machine. A new employee could receive guidance inside the application where the work occurs.

This does not eliminate formal development. It changes what formal development should accomplish.

05 LEARNING NEEDS TO MOVE CLOSER TO WORK

Formal learning remains particularly useful for:

foundational knowledge

mental models

judgment

deliberate practice

complex skills

reflection

social learning

experiences where failure must occur safely before it occurs at work

Performance support handles more of what people previously had to remember.

The boundary between learning and doing becomes less important. Employees were never hired to demonstrate learning. They were hired to perform.

L&D needs a different response to speed

AI's most visible contribution to learning is speed. An instructional designer can generate a draft in minutes. Scenarios that once took days can be produced almost instantly. Assessments, videos, simulations, translations and job aids can all be created faster.

The obvious concern is quality. That concern is justified. Poor thinking can now produce polished-looking training remarkably quickly.

But deliberately slowing production misses the larger opportunity. The better use of speed is experimentation.

Suppose a business unit has a performance problem. The traditional development cycle might involve weeks of analysis, several months of design and production, a polished launch, and an evaluation conducted after substantial resources have already been committed.

EXHIBIT Use speed to shorten the evidence cycle

Old emphasis	Better use of AI speed
Idea → finished course	Idea → small intervention
Polish before exposure	Expose to real users early
Launch at scale	Test with limited scope
Evaluate after investment	Collect evidence before expansion
Defend sunk cost	Scale, adapt or stop

AI should shorten the distance between an idea and evidence.

The minimum effective intervention

L&D has a natural tendency toward completeness. If a problem is important, we build a program. If a capability matters, we create a curriculum. If the audience is large, we build scalable content.

There is another approach: identify the smallest intervention capable of producing the required change. Call it the Minimum Effective Intervention.

It might be a checklist, a manager conversation, three practice scenarios, an AI prompt, a workflow change, a short demonstration, a performance aid, targeted coaching, a five-minute explanation, a redesigned form or removal of an unnecessary approval step.

The purpose is not to make learning smaller for its own sake. It is to avoid solving more problem than actually exists.

The minimum intervention also creates a faster feedback loop. If it works, expand it. If it partially works, improve it. If it does nothing, stop. The organization learns before it spends heavily.

EXHIBIT Minimum Effective Intervention decision path

1	2	3	4	5
Define the performance gap	Identify the real constraint	Choose the smallest plausible intervention	Test in work	Scale only on evidence

How much intervention does this problem genuinely deserve?

Replace the learning portfolio with a capability portfolio

Many L&D teams manage portfolios of programs: leadership development, compliance, onboarding, sales training, professional development and technical learning.

There is nothing inherently wrong with that structure, but it makes learning activity the organizing principle. A capability portfolio starts with business performance.

For every significant investment, L&D should be able to answer seven questions.

1	What outcome matters? Revenue, quality, productivity, customer experience, safety, speed, retention or another material result.
2	What must people do differently? Describe observable behavior, not what people should "understand."
3	Why is that behavior not happening already? Identify the actual constraint.
4	What is the minimum intervention likely to change it? Do not assume training.
5	What role should AI play? Automate, assist, coach, retrieve, analyze, generate or no role at all.
6	What must remain human? Identify where judgment, accountability, expertise and relationships matter.
7	What evidence will tell us it worked? Define this before development begins.

08 REPLACE THE LEARNING PORTFOLIO WITH A CAPABILITY PORTFOLIO

The portfolio can then be ranked by expected value.

This has an important consequence. Some learning requests should be declined. Others should be reduced. A small number should receive substantially more attention because the underlying capability has large business consequences.

Google describes something similar in its approach to enterprise AI, recommending a focused portfolio of roughly five to seven high-impact use cases rather than allowing scattered experimentation to consume resources.

L&D could benefit from the same discipline.

Portfolio rule

Fewer capability problems. Deeper diagnosis. Better evidence. More deliberate scaling.

Measure the human-AI system, not the technology in isolation

Organizations routinely measure whether AI can perform a task. A more useful question is whether a person performs better with AI.

Stanford researchers call this a "centaur evaluation": testing the performance of humans and AI working together rather than treating either one as the isolated unit of analysis.

That idea is particularly useful for L&D.

Suppose an organization gives managers an AI tool to help prepare for performance conversations. The wrong evaluation asks: How good are the AI's recommendations? The stronger evaluation asks: Do managers using the tool conduct better conversations?

Likewise: Do salespeople with AI research support prepare faster and ask better questions? Do new service agents with AI assistance resolve more cases without reducing quality? Do instructional designers working with AI create effective interventions faster?

EXHIBIT The unit of performance

Person	+	Technology	+	Workflow
Judgment, accountability, skill		AI assistance, automation, retrieval		Process, tools, constraints, reinforcement

Training metrics may still be useful diagnostically. They should not be confused with the result.

- cycle time
- quality
- error rates
- customer outcomes
- productivity
- consistency
- escalation rates
- time to proficiency
- decision quality
- rework
- employee effort
- human judgment required

AI literacy is necessary, but broad AI training is not a strategy

The Stanford AI Index provides useful context for the urgency. Organizational AI adoption reached 88% in the data reported in the 2026 Index, while generative AI adoption has spread at a pace faster than previous general-purpose technologies such as the personal computer and internet. The report also finds a persistent gap between adoption and governance, with responsible AI practices struggling to keep pace with technical capability.

Some baseline AI capability will therefore become ordinary workplace literacy.

Employees need to understand what the technology can do, where it fails, how to verify outputs, what information can be shared, when human review is required, how to use approved tools, where accountability remains and when not to use AI.

But an organization should be cautious about turning this into a massive generic training program.

A marketer, claims adjuster, software engineer and frontline manager do not need identical AI capability.

After a small common foundation, development should move quickly into work. Teach AI through the tasks employees actually perform. Give them relevant examples. Let them practice on realistic problems. Show where AI improves the work and where it introduces risk. Measure whether the resulting workflow performs better.

AI literacy matters. Applied AI capability matters more.

Governance should appear where decisions are made

Responsible AI is often treated as a policy problem. Legal develops requirements. Technology teams implement controls. Employees complete mandatory training. Then everyone returns to work.

That model is weak because most AI risk appears during ordinary decisions.

Can this customer information be entered into the system? Should this recommendation be trusted? Does this output need human review? Can AI generate this communication? What happens when the system produces conflicting information?

The Stanford AI Index reports that responsible AI measurement and reporting continue to lag behind capability development while documented AI incidents have increased.

L&D has a role here, but it should resist turning governance into another annual awareness module.

EXHIBIT Governance as an operating system

Layer	Purpose
Policy	Defines what is permitted and prohibited.
Technology	Prevents, flags or constrains dangerous behavior.
Workflow	Places prompts and guidance where risky decisions occur.
Capability	Builds judgment for situations rules cannot answer.

Training becomes part of the control environment rather than the entire control environment.

The instructional designer becomes a capability designer

Generative AI will almost certainly change instructional design. The least interesting version of that change is faster production.

AI can already assist with first drafts, content transformation, basic scenarios, assessment items, media production, summaries, localization and research support.

As those tasks require less human time, the question becomes what instructional designers should do with the capacity.

The answer should not be "produce even more content." Move the effort to work that requires stronger judgment.

EXHIBIT Where instructional design value moves

Less production effort	More capability work
First drafts	Diagnosing performance
Basic content transformation	Understanding workflows
Routine assessment items	Observing users
Simple scenarios and media	Identifying behavioral barriers
Formatting and localization	Designing practice and support
Content assembly	Testing interventions and interpreting evidence

12 THE INSTRUCTIONAL DESIGNER BECOMES A CAPABILITY DESIGNER

A capability designer is accountable for helping create performance, not merely learning.

Sometimes the answer will still be a course. Sometimes it will be a practice environment. Sometimes an AI coach. Sometimes a checklist. Sometimes manager reinforcement. Sometimes removing a barrier that made training unnecessary in the first place.

That makes the work harder in one sense. It also makes it considerably more valuable.

The deliverable is whatever combination of people, process, technology and learning best creates the required capability.

L&D needs its own AI operating model

The traditional ADDIE sequence assumes a reasonably stable problem that can be analyzed, designed for, developed, implemented and evaluated. That logic remains useful in some settings. It is too slow and linear for many rapidly changing capability problems.

A more suitable operating model is cyclical.

EXHIBIT The capability loop

Sense	Diagnose	Prioritize	Design	Test	Measure	Adapt	Scale
Detect change	Find the real constraint	Focus on business value	Choose the minimum solution	Expose early	Observe performance	Improve or stop	Expand what earned it

13 L&D NEEDS ITS OWN AI OPERATING MODEL

EXHIBIT Operating rhythm

Step	Question
Sense	Continuously identify changes in strategy, technology, work, customer expectations and performance. Do not wait for training requests. Look for emerging capability gaps.
Diagnose	Determine what people need to do and what prevents them from doing it. Separate knowledge problems from environmental problems.
Prioritize	Estimate the consequence of solving the problem. Focus resources where performance matters.
Design	Select the smallest combination of learning, technology, process, management action and performance support likely to work.
Test	Expose the intervention to real users early. Capture failures before scale magnifies them.
Measure	Look for changes in behavior and performance. Where possible, compare against a reasonable baseline or control.
Adapt	Improve what works. Change what partly works. Stop what does not.
Scale	Only after the intervention has earned the right to scale should L&D invest heavily in reach, standardization and infrastructure.

13 L&D NEEDS ITS OWN AI OPERATING MODEL

Measurement discipline matters

Stanford's Enterprise AI Blueprint makes an important point here. Observed improvement can easily be mistaken for causal impact.

In one example, users of an AI coding tool initially appeared more than 50% more productive. After accounting for the fact that the strongest developers had adopted the tool first, the estimated improvement was closer to 10% to 15%.

L&D should take the same warning seriously.

The fact that high performers completed a program does not mean the program created high performance.

This creates a capability loop rather than a course-development pipeline.

Measurement rule

Define the evidence before development begins. Distinguish participation, correlation and plausible causal contribution.

A practical maturity model

Organizations will not make this transition at once. A useful maturity model has four stages.

EXHIBIT L&D AI maturity

Stage	Primary question	Typical state
1. AI-assisted L&D	How can AI make L&D more efficient?	The existing model remains intact. AI accelerates drafting, research, graphics, assessments and content conversion.
2. AI-enabled learning	How can AI improve learning?	AI changes the learning experience through coaching, simulation, personalization and adaptive support.
3. AI-enabled performance	How can people perform better with AI?	AI moves into the workflow. Task analysis, performance support and business measures become central.
4. Adaptive capability system	How quickly can the organization build the capability its strategy requires?	Learning, work design, AI, expertise and evidence operate as one continuous capability system.

This is the destination, not a maturity contest. Not every organization needs to reach Stage 4 immediately. Every L&D leader should understand where the function is heading.

What L&D should do now

The transition does not require a large transformation program. It can begin with a small number of operating changes.

EXHIBIT Seven moves to start now

Move	What it means
Pick five capability problems that matter	Do not begin with AI use cases. Begin with business problems where human performance has material consequences.
Map the work	Break the relevant roles into tasks. Identify which tasks remain human, which can be assisted, which can be delegated and which should disappear.
Establish a minimum capability baseline	Define what employees must know and be able to do without assistance. Do not confuse access to AI with competence.
Move support into work	Where remembering is unnecessary, stop forcing memorization. Provide reliable support when the decision occurs.
Capture expertise	Identify high performers whose tacit knowledge would improve the wider system. Document reasoning, not just procedures.
Experiment before scaling	Build the smallest useful intervention. Test it. Measure actual performance. Then decide whether it deserves additional investment.
Change the scorecard	Keep completion data where it is operationally useful. Add measures that show whether the work improved.

The leadership question

The Stanford Enterprise AI Playbook reaches a conclusion that should make L&D leaders uncomfortable in a productive way.

Across 51 enterprise deployments, similar technology produced very different outcomes because organizations differed in readiness, leadership, processes and capacity to change.

AI transformation is therefore partly a capability problem.

That does not mean L&D should claim ownership of AI transformation. It means L&D should decide whether it intends to be relevant to it.

If the function remains primarily responsible for courses, content and completion, AI will make that work cheaper and easier to distribute.

If L&D instead helps the organization diagnose performance, redesign work, capture expertise, develop judgment, support adoption and determine whether people and AI are actually producing better outcomes, its strategic value increases.

That choice is larger than technology. It is a choice about what L&D believes its job is.

The leadership decision is whether L&D will remain a learning producer or become part of the organization's capability system.

CONCLUSION

The opportunity is to redesign how capability gets built

AI will create plenty of new learning.

Employees will need new skills. Leaders will need new judgment. Organizations will need stronger AI literacy, governance and technical capability.

But creating all of that learning is not the biggest opportunity in front of L&D.

The bigger opportunity is to reconsider how capability gets built.

Some capability will still come from formal learning. Some will come from experience. Some will sit inside workflows. Some will be supplied by AI. Some will remain concentrated in experts. Some will emerge from people and machines working together.

The L&D function has historically concentrated on one part of that system. It now has an opportunity to work on the system itself.

How can we help the organization build capability fast enough to keep up with the work?

That is a substantially larger mandate. It is also the one the business is going to need.

SOURCE FOUNDATION

Primary sources used to build the L&D adaptation

EXHIBIT Evidence base

Institution	Source	Use in this report
MIT Sloan Management Review	GenAI: The Strategy Toolkit	Paid proprietary toolkit. Only publicly available framework descriptions informed this report.
MIT Sloan	Leading the AI-Driven Organization	Business-problem-first AI roadmap, data strategy and employee capability.
MIT Sloan	How to Set Technology Strategy in the Age of AI	Value creation, capture and delivery as a strategy lens.
Stanford Digital Economy Lab	The Enterprise AI Playbook: Lessons from 51 Successful Deployments	Enterprise implementation evidence and organizational readiness.
Stanford HAI / Economist Enterprise	The Enterprise AI Blueprint	Task decomposition, organizational redesign and complementary investment.
Stanford Institute for Human-Centered AI	2026 AI Index Report	Adoption, economics, education, governance and responsible AI trends.
Google Cloud	How to Build an Effective AI Strategy	Vision, high-value use cases, measurement and scaling.
Google Cloud	Scaling AI from Experimentation to Enterprise Reality	Focused use-case portfolio and enterprise scaling.
Google Cloud	A Guide for Leaders on Implementing Agentic Solutions	Strategic alignment, ecosystem mapping, rapid prototyping and risk.
Stanford research	Generative AI at Work; Centaur Evaluations	Productivity effects, expertise diffusion and human-AI system evaluation.



CLOSING PERSPECTIVE

AI will create plenty of new learning. But the larger opportunity is to redesign how capability gets built. Some capability will live in people, some in workflows, some in AI, some in teams, and some in the interaction between them.

**How can we help the organization
build capability fast enough to keep
up with the work?**

DR. RAVINDER TULSIANI

Enterprise Capability Executive

Workforce Capability | AI Readiness | Learning Impact

rtulsiani.com

