




STRATEGY / OCTOBER 2026

Protect the learning inside the work

How AI adoption can build tomorrow's judgement while improving today's performance



THE DECISION IN BRIEF

AI-assisted delivery should develop tomorrow’s experts as well as improve today’s output; fund practice, feedback and unfamiliar-case assessment inside the redesigned role.

- 01

Removing routine preparation can also remove occasions to learn judgement. Identify which activities build expertise before deciding which work disappears.
- 02

The evidence points to design choices, not inevitable skill loss. Some assistance transmits useful expertise; answer provision can improve immediate output while weakening unaided learning in specific settings.
- 03

For each redesigned role, fund a Learning Covenant: the judgement to retain, deliberate practice, expert feedback and evidence of transfer. Assess delivery performance and capability development separately.

THE MAYDEN VIEW


Every redesigned workflow needs a route from assisted performance to independent judgement.

INSIDE THIS PERSPECTIVE

The argument at a glance	3
01 / Identify the judgement to retain	4
02 / Build practice into work	5
Exhibit 01 / An unfamiliar case tests how the analyst reaches a view	6
03 / Observe transfer and recovery	7-8
Exhibit 02 / A Learning Covenant makes development an operating commitment	9
04 / Renew the capability commitment	10
The leadership agenda	11
Method and references	12

ONE DECISION, FOUR CONNECTED QUESTIONS

Learning belongs inside the work



INTERLOCKING FEEDBACK / CONCEPTUAL ARTWORK

01 Identify the judgement to retain

Account for the expertise the old workflow produced

02 Build practice into work

Define a Learning Covenant for each redesigned role

03 Observe transfer and recovery

Fund teaching as well as approval

04 Renew the capability commitment

Keep access and opportunity in the redesign

Every redesigned workflow needs a route from assisted performance to independent judgement.

CHAPTER 01

01

Identify the judgement to retain

Account for the expertise the old workflow produced

A firm can improve today's output while weakening the route through which tomorrow's experts develop. When AI prepares the case, the junior employee may lose the occasion to identify a contradiction, weigh evidence or defend a recommendation. The question for leaders is how that judgement will now be learned. Faster preparation does not answer it.

Preserve the learning function rather than every old task. Some repetition teaches little; other work builds the mental models needed for unfamiliar cases. Identify the judgement, the conditions in which it matters and the evidence that someone can exercise it. A generic requirement for critical thinking is too broad to guide a workflow redesign.

The risk is a gap between current productivity and future capability: experienced staff keep correcting output while fewer colleagues learn to replace them. This is a hypothesis to test, not an established consequence of AI. The same technology may make expert feedback and practice more accessible. The operating design determines which possibility the firm pursues.

The evidence argues for conditional design, not a single verdict

The OECD's analysis of job vacancies finds changing demand for skills in occupations exposed to AI, including management and business skills. ^[1] Vacancy requirements are an imperfect measure of what workers actually do, and exposure is not adoption. The useful implication is that a workforce plan cannot be reduced to training more technical specialists. Organizations should investigate how the task bundle and the judgement required of ordinary employees are changing.

A field experiment with nearly 1,000 high-school mathematics students compared different AI tutor designs. The authors report that unguarded assistance could improve practice performance while harming subsequent unaided test performance; learning safeguards largely mitigated the harm without establishing a positive unaided-learning effect. ^[2] This is educational evidence, not a measured effect on corporate professionals. It establishes a reason to test whether a workplace tool supports reasoning or simply supplies answers, without importing the study's effect size into a staffing forecast.

A 2026 randomized study of 52 developers learning an unfamiliar Python library found lower subsequent quiz performance with AI assistance, without a statistically significant improvement in completion time. ^[8] The experiment concerns initial acquisition of one skill, not long-term professional competence. Its relevance is practical: when an assisted workflow also trains future reviewers, measure what people can explain, debug and recover without assistance as well as what they finish with it.

A survey of 319 knowledge workers examined self-reported changes in critical-thinking effort when using generative AI. ^[3] Such reports can reveal perceived practices and confidence, but cannot demonstrate a causal decline in ability. Treat the study as a prompt for observation: what do employees verify, what do they accept and what can they explain independently? The company should measure the capability it needs rather than infer it from either tool enthusiasm or anxiety about dependence.

CHAPTER 02

02

Build practice into work

Define a Learning Covenant for each redesigned role

mayden's Learning Covenant assigns development responsibilities alongside delivery responsibilities. The operating leader names the judgement the role must retain; the employee receives relevant practice and feedback; a competent reviewer assesses readiness for greater responsibility. It is a proposed role-design agreement, not a validated score. Its purpose is to put learning in the operating plan.

For an illustrative commercial analyst, the covenant might require the ability to distinguish a source fact from an assumption, test whether a comparison is valid and explain a recommendation's sensitivity to uncertainty. AI can help retrieve evidence and suggest alternatives. The analyst still needs occasions to form a view before seeing the proposed answer and to defend it against a meaningful challenge. The assessment should examine reasoning on a new case rather than recall of a familiar template.

For a service role, the retained judgement may concern when the standard remedy is inappropriate and how to recognize a vulnerable customer or conflicting evidence. Practice can use carefully prepared cases with known resolution criteria. A senior reviewer should explain the decisive feature, not merely replace the answer. The covenant makes the learning obligation visible in the operating plan, including the time and expertise needed to fulfil it.

Use AI to increase the quality of practice

AI assistance can create useful learning opportunities when the design makes thinking visible. A tool can ask an employee to identify missing evidence, compare alternative explanations or predict the consequence of an action before offering feedback. These are proposed design patterns requiring validation. An apparently helpful explanation can still be wrong, so the practice environment needs trusted reference material and a way for a competent reviewer to examine the feedback.

The support-agent field study reports evidence consistent with worker learning alongside productivity gains. ^[4] It provides evidence against assuming that assistance inevitably reduces skills. The management question is which features of a deployment transmit useful expertise and which merely make users dependent on an answer. That distinction should be investigated through retained performance, transfer to unfamiliar cases and the quality of escalation, not by counting completed learning modules.

Concentrate practice where it supports future responsibility and recovery. Repeating every automated task would consume the capacity the tool releases. Instead, select unfamiliar cases, define the assistance boundary and explain the judgement each exercise develops. Keep routine assistance available where unaided repetition offers little learning value.

ILLUSTRATIVE DESIGN · NO MEASURED RESULTS

An unfamiliar case tests how the analyst reaches a view

Illustrative decision record

Assessment field	Illustrative analyst task	Reviewer evidence
Case and purpose	New comparison with one unsupported assumption	Domain owner validates task and expected reasoning
Assistance boundary	Initial view without generated answer; then permitted assistance	Record tool access and changes to the initial view
Diagnosis	Separate fact, assumption and inference	Identify decisive source and missing evidence
Error and calibration	Challenge an attractive but unsupported conclusion	Explain uncertainty and whether evidence changes confidence
Recovery and feedback	Resolve without the tool, or escalate safely	Observe recovery and explain the decisive correction
Follow-up	A different case after supported practice	Compare several observations, including accessibility effects

THE MANAGEMENT IMPLICATION

A polished assisted answer does not establish diagnosis, transfer or recovery.

Proposed developmental card, not a validated employment score, psychometric scale or universal pass mark. Use paid practice, competent review and multiple observations; explain use and challenge rights. Source: mayden analysis. Proposed framework or illustrative scenario.

CHAPTER 03

03

Observe transfer and recovery

Fund teaching as well as approval

A senior professional's approval and teaching serve different purposes. Approval releases a result; teaching develops someone able to reach and challenge a result later. Separate the two in workload and performance expectations.

Otherwise the immediate queue can absorb all senior attention while the firm's next generation of reviewers remains underdeveloped.

Capture concise explanations of consequential decisions in an authorized knowledge record. Include the context, the evidence that changed the answer and the limits of the reasoning. Avoid presenting a past judgement as a universal rule. A useful case library preserves disagreement and changing conditions as well as successful

resolution. Employees should be able to see why superficially similar cases led to different outcomes, which is often where professional judgement becomes visible.

Protect time for this work. If senior staff are measured only on immediate delivery volume, teaching will compete with an urgent queue it rarely wins. The OECD's 2026 skills brief emphasizes that training belongs within a broader response to workplace AI. ^[5] mayden's application is to place capability development inside the operating budget and leadership review. A learning commitment without time, feedback or consequences for neglect is unlikely to survive production pressure.

CHAPTER 03

03

Observe transfer and recovery

Measure transfer, calibration and recovery

Assisted performance answers whether the employee and system can produce an acceptable result together. Transfer asks whether the employee can apply the relevant judgement to an unfamiliar case. Calibration asks whether the person recognizes the limits of their own and the system's knowledge. Recovery asks whether they can act when the tool is unavailable or wrong. These are distinct assessment questions. Strong performance on a familiar assisted task should not silently answer all of them.

Exhibit 01 provides a developmental record for an unfamiliar analyst case. The domain owner sets the expected reasoning and assistance boundary. The employee forms a view, then explains what permitted assistance changes. A reviewer records diagnosis, uncertainty, response to a misleading answer and recovery, and gives a reasoned correction. Repeat with a different case after practice. Explain the retained evidence and appeal route: one exercise cannot establish employment fitness, and language or accessibility barriers must not be mistaken for weak domain judgement.

Track development at the role and cohort level as well as individual performance. Are people becoming able to resolve more difficult cases? Is the same small group still rescuing every exception? Are promotions supported by demonstrated judgement? Use multiple observations over time, because a single exercise can reflect familiarity, language or test design. The goal is a credible account of capability development, not a new universal score that gives false precision to a complex learning process.

Promotion criteria should evolve with the redesigned role. An employee who produces polished assisted output may be ready for more responsibility, but the assessment should also examine whether they can frame the problem, identify missing evidence and challenge an attractive wrong answer. Preserve examples of the reasoning behind these judgements so that development conversations are concrete. Where an employee lacks capability, distinguish insufficient practice from an unsuitable assessment or inaccessible tool. The response may be a better learning opportunity rather than a negative performance label. This makes the capability account useful to both the organization and the person whose work is changing.

ORIGINAL MAYDEN FRAMEWORK · ILLUSTRATIVE ROLE DESIGN

A Learning Covenant makes development an operating commitment

Illustrative role design



THE MANAGEMENT IMPLICATION

Learning needs time, feedback and accountable ownership within the workflow.

Original mayden synthesis; no claim of measured training effectiveness. Source: mayden analysis. Proposed framework or illustrative scenario.

CHAPTER 04

04

Renew the capability commitment

Keep access and opportunity in the redesign

Exposure estimates from the ILO indicate that many occupations contain tasks potentially affected by generative AI; they do not predict individual job losses. ^[6] Workforce planning should therefore examine how opportunities to learn and advance change within roles. If some employees receive better tools, protected practice and senior attention while others receive only higher output targets, the implementation may widen capability differences even when access to the basic software is nominally equal.

Consider language, accessibility and prior experience in both tools and assessment. A person may understand the domain while struggling with an interface or a training exercise written in a second language. Conversely, fluent interaction with the tool may conceal weak domain judgement. Separate these issues before deciding what development is needed. The organization should provide routes to demonstrate competence that resemble the work and do not reward presentation polish as a substitute for understanding.

UNESCO's guidance advocates human agency and pedagogical validation in educational uses of generative AI. ^[7] Applied to workplace development, this supports an explicit review of who benefits from the learning design and whether it actually teaches the intended capability. It does not make an enterprise programme educationally validated. Leaders should test participation, feedback quality and transfer, then adjust the design when a group is systematically excluded or when apparent progress depends on continued answer provision.

Invest in the next generation of experts

Technology makes some skills obsolete; retaining them all would be wasteful. Preserve judgement needed to frame work, assess evidence, recognize failure and resolve current exceptions. Retire practice that no longer serves those responsibilities, and invest in the capabilities the redesigned role now requires.

Future systems may handle more exceptions. Reassess the covenant as reliability and responsibility change, using evidence rather than a forecast of autonomy. Human practice should support a defined need; equally, an unproven future capability should not justify removing a present recovery skill.

Require each major workflow redesign to explain how its next generation of experts will develop. Approve the retained judgement, practice time and feedback capacity alongside the delivery plan. Then review unfamiliar-case performance and assisted output separately. A stronger result today should not conceal an unfunded capability requirement for tomorrow.

FROM INSIGHT TO ACTION

Decisions to take forward

OPERATING LEADER

Fund the learning covenant

Evidence to require: Retained judgement and paid practice capacity

DOMAIN REVIEWER

Select feedback and next practice

Evidence to require: Reasoning, error recognition and repeated transfer cases

PEOPLE LEADER

Prevent unsupported employment conclusions

Evidence to require: Access, assessment purpose and challenge route

mayden connects senior judgement, research and engineering to turn strategic choices into working capabilities.

EVIDENCE YOU CAN INSPECT

Research notes

Method and boundaries

Narrative synthesis of eight primary research and institutional sources, with an original mayden role-design proposal. No pooled effect estimate or proprietary workforce study.

Contrasting evidence on learning, perceived cognitive effort, task exposure and skills policy; all sources opened 10 October 2026.

Education results do not establish workplace effect sizes. Survey perceptions do not prove causal skill decline. The Learning Covenant requires local testing and is not a validated assessment instrument.

Primary sources and author research

[1] Artificial intelligence and the changing demand for skills in the labour market

Andrew Green; OECD / 10 April 2024. Abstract; executive summary. Vacancies are not observed job performance; exposure differs from adoption.

[Read the source](#)

[2] Generative AI Without Guardrails Can Harm Learning: Evidence from High School Mathematics

Bastani et al.; author manuscript for PNAS / 2025. Author manuscript pp. 1-6; abstract, experimental design and footnote 3. One Turkish high-school mathematics setting; short-term unaided exam performance, not long-term retention or workplace productivity. Safeguards largely mitigated harm without demonstrating positive unaided-learning effects.

[Read the source](#)

[3] The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers

Lee et al.; Microsoft Research and Carnegie Mellon; CHI2025 / 2025. Abstract; methods; discussion of self-report limitations. Not a causal test of cognitive ability or longitudinal skill loss; vendor-affiliated research.

[Read the source](#)

[4] Generative AI at Work

Brynjolfsson, Li and Raymond; arXiv v 2 / 6 November 2024. Abstract; learning findings. One firm and workflow; does not establish that every assistance design develops judgement.

[Read the source](#)

[5] AI and skills: What we know so far

OECD / 5 June 2026. Training and broader policy response. Policy guidance, not a causal evaluation of the proposed workplace covenant.

[Read the source](#)

[6] Generative AI and Jobs: A Refined Global Index of Occupational Exposure

International Labour Organization / 20 May 2025. Overview; task-exposure methodology. Not a forecast of individual employment or skill loss.

[Read the source](#)

[7] Guidance for generative AI in education and research

UNESCO; Miao and Holmes / 7 September 2023; webpage updated 16 January 2026. Human-centred approach; pedagogical validation. Educational guidance used by analogy; not evidence that a workplace programme is effective.

[Read the source](#)

[8] How AI Impacts Skill Formation

Shen and Tamkin; arXiv / 28 January 2026, v1. Sections 5.2.1-5.2.2; Figure 6; limitations discussion. Small short-session sample; no long-term workplace transfer measured. Interaction-pattern groups are observational, not separately randomized.

[Read the source](#)