


STRATEGY / OCTOBER 2026

An AI-native enterprise starts with the work

Why the unit of transformation is the complete customer outcome



THE DECISION IN BRIEF

Redesign the constraint that prevents a completed outcome. AI creates enterprise value when the surrounding process can convert better task performance into accepted work.

- 01

Start with one outcome somebody accepts: a resolved case, a usable supplier record or an implemented decision. Trace what prevents it from being completed.
- 02

Choose among process repair, AI assistance and sequence redesign according to that constraint. Task-level gains support a testable opportunity; they do not establish an enterprise-wide automation return.
- 03

Fund the smallest complete change that improves the outcome. Measure displaced review, exceptions and how released capacity will be used before scaling it.

THE MAYDEN VIEW


A faster task creates value only when the organisation can turn it into completed work.

INSIDE THIS PERSPECTIVE

The argument at a glance	3
01 / Locate the unfinished outcome	4-5
02 / Choose the intervention	6
Exhibit 01 / One supplier outcome, three competing interventions	7
03 / Convert task gains into accepted work	8-9
Exhibit 02 / Exposure is broader than the highest-exposure category	10
04 / Fund the smallest complete change	11
The leadership agenda	12
Method and references	13

ONE DECISION, FOUR CONNECTED QUESTIONS

Work becomes a system



ACCOUNTABLE WORK GRAPH / CONCEPTUAL ARTWORK

01 **Locate the unfinished outcome**
Find the constraint behind the unfinished outcome

02 **Choose the intervention**
Choose process repair, assistance or redesign

03 **Convert task gains into accepted work**
Measure accepted work and the use of released capacity

04 **Fund the smallest complete change**
Test the simpler process before a wider transformation

A faster task creates value only when the organisation can turn it into completed work.

CHAPTER 01

01

Locate the unfinished outcome

Find the constraint behind the unfinished outcome

A team drafts twice as many responses but resolves no more customer problems. The constraint may be refund authority, specialist availability or a broken customer record. More generation then enlarges the queue. Before selecting a tool, name the outcome that must improve and locate the condition preventing it. That is the starting point for an AI-native enterprise.

Consider an illustrative service journey from request to closure. Reading the request, retrieving evidence, deciding the remedy, obtaining approval and confirming resolution are connected activities. Improving the first two may help, but a customer experiences their combined result. A redesign might eliminate unnecessary

approval for low-value, reversible remedies while escalating ambiguous cases with the supporting record attached. The AI contribution is useful because the surrounding process permits it to change an outcome. This is an analytical example, not a mayden client result.

The economic logic has a formal precedent. Agrawal, Gans and Goldfarb model adoption across interacting tasks and show why coordination and system structure matter. Their work supports a design question, rather than a promise that broader transformation always wins. ^[1] mayden's interpretation is that leaders should examine the dependency around a task before estimating its value. A narrow assistant can be the right answer when the task is already independent; an interconnected workflow needs a different investment case.

CHAPTER 01

01

Locate the unfinished outcome

Transfer the mechanism, not the headline gain

A field study covering 5,172 customer-support agents found a 15% average increase in issues resolved per hour, with heterogeneous effects by worker experience. ^[2] The outcome is operationally meaningful because it includes resolution, rather than merely generated text. It is still evidence from a particular company and deployment. An executive should borrow the study's discipline of measurement and segmentation, then test whether the same mechanism exists in the organisation's own service work.

A randomized writing experiment with 453 professionals reported 40% less time and 18% higher assessed quality on selected assignments. ^[5] These are strong task-level findings, but the assignments do not represent the full burden of organizational work. Company-specific facts, stakeholder alignment, legal review and implementation can dominate elapsed time. Before importing an effect into a business case, specify which portion of the actual workflow resembles the experimental task and which portion does not.

Occupational exposure is an equally useful but different lens. The ILO's 2025 assessment puts one quarter of employment in occupations with some generative-AI exposure, while the highest category covers 3.3% globally. ^[3] Exposure should prompt examination of task bundles. It should not be converted into a headcount target. The important strategic possibility is that removing a routine step changes what the remaining job can accomplish, including service quality or coverage that was previously unaffordable.

Task exposure is an incomplete guide to where adoption will occur. A 2026 Federal Reserve research summary describes nationally representative survey evidence showing substantial differences between people doing similar work; technical exposure explains only part of observed use. ^[8] For mayden, this strengthens the case for mapping the actual workflow with its users. Access, incentives, tacit knowledge and the consequences of error are hypotheses to investigate locally, rather than reasons to assume that every exposed task should be automated.

CHAPTER 02

02

Choose the intervention

Choose process repair, assistance or redesign

mayden's Outcome-Constraint-Authority map puts three questions on one page. What does the recipient accept as complete? What limits completion? Who can change that condition? The constraint may be expertise, evidence, permission, capacity or demand. This proposed framework is a diagnostic, not a validated maturity score. A tool idea becomes an executable case only when all three questions have answers.

Use routine supplier onboarding as a bounded example. Completion means an approved, usable record with required evidence and no unresolved control issue. Exhibit 01 compares process repair, AI assistance and redesign against that same result. If requests wait for an undefined approval threshold, first test a policy repair; faster document drafting does not grant authority. If repeated evidence assembly instead consumes specialist time, test assistance while preserving specialist acceptance. Redesign the sequence only when a case trace shows that handoffs, rather than a single task, determine delay. These are conditional choices, not observed rankings. Procurement must establish where time goes before funding the intervention. After preparation improves, repeat the trace: a growing reviewer queue can become the next binding constraint, invalidating a scale-up based only on faster preparation.

This mapping reveals when the supposed AI problem is a policy problem. If every request needs senior approval because thresholds were never specified, a model cannot invent legitimate authority. If the same evidence is requested repeatedly because systems cannot share a record, better prose adds little. Conversely, if experts repeatedly reconstruct information from dispersed sources, retrieval and synthesis may relieve the real bottleneck.

The map converts general enthusiasm into a testable explanation of how the work could improve.

Give exceptions a destination and an owner

Straightforward cases make demonstrations persuasive. Exceptions determine whether a redesigned process can operate. A system must recognize missing evidence, contradictory records and requests outside its remit. It also needs a destination for those cases, an expected response time and a way to resume work after a decision. An escalation queue without capacity or ownership is merely a hidden manual department. Its cost and delay belong inside the original design.

Risk frameworks help articulate this boundary. NIST's AI RMF connects context, measurement and management with governance; it is voluntary guidance, not proof of operational safety. ^[6] mayden's application is to define a release boundary in ordinary business language: which cases may pass, which must be reviewed and which must stop. The boundary should reflect consequence and reversibility. A reversible draft and an irreversible payment instruction should not inherit the same review path merely because the same model produces both.

Record the conditions under which the system was tested. Model-card research emphasizes intended uses and performance in relevant conditions. ^[7] For an enterprise workflow, the analogous record should include language, customer segment, case complexity and evidence availability. Average quality can conceal a failure affecting a smaller but consequential group. Such documentation makes a handover practical: an operational owner can understand what the system is allowed to do without reconstructing the development team's assumptions.

ILLUSTRATIVE DESIGN · NO MEASURED RESULTS

One supplier outcome, three competing interventions

Illustrative decision record

Alternative	Constraint hypothesis	Evidence needed	Conditional decision
Repair the process	Undefined approval threshold causes waiting	Trace waiting separately from evidence work; control owner confirms a lawful threshold	Repair the rule first if waiting dominates
Assist evidence assembly	Specialists repeatedly collect the same material	Compare complete accepted records, total effort and exceptions	Use assistance if accepted work improves without displaced review burden
Redesign the sequence	Repeated handoffs limit completion	Compare handoff delay, control coverage and transition burden	Redesign only if coordination gains justify the commitment

THE MANAGEMENT IMPLICATION

Recheck the next constraint: faster preparation may transfer the bottleneck to reviewer capacity.

Fictional supplier workflow. Same completion definition and observation window for every alternative; no measured ranking or promised saving. Procurement and control owners retain authority. Source: mayden analysis. Proposed framework or illustrative scenario.

CHAPTER 03

Convert task gains into accepted work

03

Measure accepted work and the use of released capacity

The primary measure should follow the work to acceptance. For customer support, that may mean durable resolution after a defined reopening window. For procurement, it may mean a usable approval with no missing control evidence. For engineering, it may mean a change that survives review and normal operation. Counting drafts, sessions or generated tokens describes activity. Those measures can explain behaviour, but they cannot establish that the enterprise is better off.

Pair a primary outcome with guardrails and a conversion account. The outcome measures throughput, quality or elapsed time. Guardrails cover unacceptable errors, complaints and displaced effort. The conversion account identifies what happens to released capacity: fewer overtime hours, shorter waiting lists, more complex cases handled or lower cash expenditure. Reallocated hours have value, but they are not automatically savings. Finance and operations should agree this distinction before a pilot creates expectations that cannot be fulfilled.

The US Census Bureau observes that AI-using firms also report training, workflow changes and cloud investment. ^[4] That association is consistent with complementary change, but it cannot establish that adoption caused better performance. A company should therefore build its own comparison. Where feasible, randomize access or phase implementation across comparable teams. Where that is impractical, document changes in demand, staffing and case mix so that a seasonal improvement is not accidentally credited to the tool.

A useful measurement boundary also has a time dimension. A response accepted immediately may be reopened after the customer tries the remedy. An approved record may fail when another team first uses it. Choose an observation window that reflects the outcome, and report provisional results separately from mature ones. Otherwise a faster process can appear better simply because its defects have not yet had time to surface. The choice should be made before results are compared, with the same window applied to the current and redesigned process.

CHAPTER 03

Convert task gains into accepted work

03

Preserve the work that develops judgement

Redesign also redistributes opportunities to develop judgement. If experienced people handle every ambiguous case while juniors only approve machine drafts, the firm may improve current throughput while narrowing its future expert pipeline. The workflow map should therefore mark tasks that teach diagnosis, evidence assessment and customer judgement. These tasks can be assisted, but their learning function should remain explicit. This is a design consideration, not a claim that every form of AI assistance reduces skill.

Use escalation as a teaching mechanism. A senior decision should return with the decisive evidence and a concise explanation of why the case crossed the boundary. Periodically ask less experienced staff to reach a view before

revealing the machine's recommendation. Track whether they become more capable of resolving comparable cases with less support. A service system that transfers judgement responsibly can expand capability; one that merely transfers approvals can concentrate dependence in a few specialists.


Worker involvement also improves the description of actual work. Process diagrams written from policy documents often miss unofficial coordination and local workarounds. Staff can identify where information is repeatedly re-entered, which exceptions consume attention and which controls are meaningful. Leaders should distinguish observations from preferences and test proposed changes, but excluding the people who perform the work removes a direct source of operational evidence from the design process.

EMPIRICAL ESTIMATE · OCCUPATIONAL EXPOSURE


Exposure is broader than the highest-exposure category

% of global employment

Some GenAI exposure



Highest exposure category



The highest-exposure category is contained within the broader 25%.

THE MANAGEMENT IMPLICATION

Use broad exposure to locate work for assessment, then examine the actual task bundle.

ILO 2025 occupational exposure model. The highest category is a subset of some exposure; bars must not be added. One-in-four represented as 25%. These are exposure estimates, not job-loss forecasts. Sources: [3]

CHAPTER 04

04

Fund the smallest complete change

Test the simpler process before a wider transformation

System redesign can become an expensive response to a modest opportunity. A small business with a simple workflow may gain more from a secure off-the-shelf assistant than from a transformation programme. A regulated process may contain deliberate separation of duties that should survive even when technology makes consolidation easy. A process with unstable demand may not support a reliable estimate of benefit. The ambition should match the economic and operational problem rather than the vocabulary of an AI strategy.

Another objection is that models improve too quickly for careful redesign. The opposite can also be true: a rapidly changing component makes durable interfaces and explicit authority more valuable. Still, a firm should avoid hard-coding unnecessary dependencies around today's limitations. Separate enduring requirements, such as evidence ownership and customer remedies, from replaceable implementation choices. Test substitution costs before committing the workflow to one supplier's proprietary structure.

Compare the proposed AI intervention with a simpler process repair: remove an unnecessary approval, clarify a form or improve access to existing knowledge. If that captures most of the benefit with less maintenance, choose it. Preserve management attention for work where AI materially changes the feasible service or operating model.

Fund the smallest complete change

Choose an outcome, an executive who can own it across functions and the constraint to test first. Agree how any released capacity will be used. These decisions establish the investment mechanism: which work improves, why it changes and where the benefit appears. A target number of pilots leaves all three unresolved.

The first release should complete a bounded piece of real work using the intended sources, reviewer and fallback. Include difficult cases and a common observation window for the existing and redesigned process. Record implementation effort as well as ongoing effort. Fund the next stage only when the complete result improves without an unowned burden moving elsewhere.

Repeat the diagnosis when scaling. A second team may face a different constraint, and faster preparation may make review the new bottleneck. Reuse the outcome definition, authority rules and evidence where they transfer; redesign what does not. AI-native operation is a discipline of changing work around demonstrated constraints, rather than distributing the same interface everywhere.

FROM INSIGHT TO ACTION

Decisions to take forward

CEO

Choose the bounded workflow to fund

Evidence to require: Common outcome and credible simpler alternative

PROCUREMENT AND CONTROLS

Select repair, assistance or redesign

Evidence to require: Waiting, evidence work and exception trace

OPERATING OWNER

Expand only after testing the next constraint

Evidence to require: Accepted throughput and review queue

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

EVIDENCE YOU CAN INSPECT

Research notes

Method and boundaries

Structured narrative synthesis of eight primary or author sources, combined with original mayden workflow analysis. No statistical pooling, proprietary fieldwork or company maturity scoring.

Selected causal and field productivity evidence, a formal organizational model, official exposure/adoption research and deployment frameworks. Sources accessed 10 October 2026.

Results span different tasks, dates and countries. The proposed framework is an analytical contribution requiring local validation; it is not an empirically calibrated diagnostic. Numerical findings retain their original populations.

Primary sources and author research

[1] AI Adoption and System-Wide Change

Agrawal, Gans and Goldfarb; University of Toronto author paper / May 2021. Abstract and model, pp. 1-16. Not an empirical estimate of the return to workflow redesign.

[Read the source](#)

[2] Generative AI at Work

Brynjolfsson, Li and Raymond; author manuscript, arXiv v2 / 6 November 2024. Abstract. One company and task setting; different older versions report different sample sizes.

[Read the source](#)

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

International Labour Organization / 20 May 2025. Overview, exposure-gradient findings. Exposure measures technical task potential, not observed replacement or a forecast.

[Read the source](#)

[4] Tracking Firm Use of AI in Real Time: A Snapshot from the Business Trends and Outlook Survey

US Census Bureau; Bonney et al. / March 2024. Abstract; CES-24-16. Associations do not identify the causal effect of adoption; data September 2023-February 2024.

[Read the source](#)

[5] Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence

Noy and Zhang; author manuscript for Science / July 2023. Abstract and experiment design. Short selected writing tasks, not whole jobs, company profit or long-term output.

[Read the source](#)

[6] Artificial Intelligence Risk Management Framework (AI RMF 1.0)

National Institute of Standards and Technology / January 2023. Part 2; GOVERN, MAP, MEASURE and MANAGE tables. Guidance is not experimental proof that any specific governance structure works.

[Read the source](#)

[7] Model Cards for Model Reporting

Mitchell et al.; Google Research / 2019. Abstract. Documentation can expose limits; it does not independently validate a deployment.

[Read the source](#)

[8] What Work Does Generative AI Do?

Bick, Blandin, Deming and Schumacher; Federal Reserve research summary / 7 July 2026. Research summary; linked St. Louis research. Self-reported adoption; no causal test of workflow redesign. July Federal Reserve version differs from August NBER record.

[Read the source](#)