


CROSS-SECTOR RESEARCH

What makes an AI-native enterprise?

A research synthesis on work, organizational capability and the evidence needed to measure them.



ORIGINAL MAYDEN GRAPHIC / OPERATING QUESTIONS

01 / THE EXECUTIVE VIEW

Define the changed work before judging the enterprise.

A research synthesis on work, organizational capability and the evidence needed to measure them.

80

selected entity records

272

recorded claims

224

qualifying claim records

9

candidate research themes

01 / THE OPERATING QUESTION

Define a workflow, responsible decision owner and evidence of repeated useful work. AI-native is Mayden's proposed organisational construct, not a status established by purchasing technology.

02 / THE OPERATING QUESTION

Keep operating practices, outcomes and supporting evidence separate. Combining them too early can reward confident disclosure rather than useful change.

03 / THE OPERATING QUESTION

Public research informs this design. A confidential assessment is in development and must pass human and empirical gates before supporting benchmark claims.

AI-assisted, bounded public-source research. Purposefully selected entity scopes; no company participation, endorsement, independent assurance or population inference implied. Confidential assessment is in development.

02 / THE ARGUMENT

Four questions to carry into the business.

Read each step as a test of the operating claim. This is an analytical framework, not a measured maturity ladder.

- 1 Which recurring work changes?**
Identify a recurring work unit, its users and the accepted result. A demonstration or purchased licence does not establish a changed operating model.
- 2 Who can act, intervene and stop?**
State the permitted action, accountable owner, intervention rights and stopping conditions. Greater autonomy is not automatically a better operating design.
- 3 What evidence establishes value?**
Keep practices, outcomes and evidence separate. Evaluate accepted work alongside quality, verification, rework, cost and an appropriate comparison.
- 4 What can a benchmark compare?**
Compare defined scopes using a tested construct and suitable evidence. Public disclosures cannot be converted into confidential participant norms.

How to use this report

Follow the source accounts and their limits, use the diagrams to frame an evidence request, then choose one recurring workflow for deeper examination. A company name identifies an attributed account; it does not imply participation, endorsement or an enterprise rating.

03 / THE ANALYSIS

Define the changed work before judging the enterprise.

Define the changed work

Mayden proposes an AI-native enterprise as one that repeatedly converts AI capability into useful, accountable work across a defined organisational scope. This working definition concerns recurring decisions, workflows, responsibilities and learning. It is a research proposition, not an empirically validated classification.

Start with a resolved service request, reviewed clinical note, accepted production schedule or authorised maintenance decision. State what AI contributes and where the work finishes. A demonstration shows that an output can be produced under particular conditions. A changed operating model needs a repeatable relationship between output, people and result.

Counterexamples sharpen the definition. Buying licences establishes availability, not active use. A specialised model can be valuable without transforming an organisation. Restricting automation can reflect sound judgement; maximum autonomy is not the objective.

Test the mechanism and work boundary

AI research does not deliver one transferable enterprise effect. A customer-support study found gains differing by experience; writing experiments and a bounded consulting exercise study other work. They identify mechanisms and questions but cannot be averaged into an expected return for a Gulf enterprise. [Generative AI at Work; Experimental evidence on the productivity effects of generative artificial intelligence; Navigating the Jagged Technological Frontier.](#)


Evidence complicating the positive story belongs in the same review. An early-2025 developer study found slower completion in its small selected setting despite perceived speed-up. A subsequent methods update describes selection and time-measurement problems in newer estimates. This warns about transfer and measurement; it is not a universal conclusion about developers. [Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity; We are Changing our Developer Productivity Experiment Design.](#)

Define eligible work, observation period, comparison and acceptance criterion before evaluating results. Include rejected outputs, checking, interruptions and the people absorbing change. A faster component can shift work elsewhere.

EXHIBIT 01 / CONCEPTUAL TOOL

A workflow claim needs four different tests

Mayden analytical framework / questions to test, not measured maturity stages



Mayden conceptual framework. These questions specify evidence to request; they do not assert that a named organisation has or lacks the practice. This is not a validated instrument or an observed causal sequence.

Unit: Conceptual decision or evidence request. Research cutoff: 2026-10-10.

03 / THE ANALYSIS

Keep three evidence ledgers

The practice ledger records workflow, scope, accountable role, permissions, frequency and coverage. The outcome ledger records quality, time, throughput, cost, service and adverse effects, together with the comparison supporting interpretation. The evidence ledger records respondent role, document or observation, date, locator, provenance, contradictions and review.

Connect the ledgers without collapsing them into one number. A documented disappointing result is scientifically valuable. A reported saving without a baseline may merit further study. Strong source access does not establish effectiveness; limited public disclosure does not establish a weak company.

Nine candidate themes organise the inquiry: AI in core work; workforce; ecosystem; interfaces and automation; decision lineage; experimentation; autonomy and operating model; organisational knowledge; governance. They are hypotheses about a useful composite, not interchangeable symptoms of a proven latent trait. Internal consistency alone would not validate them.

[Conventional Wisdom on Measurement: A Structural Equation Perspective](#); [Handbook on Constructing Composite Indicators: Methodology and User Guide](#).

Build the benchmark around a useful decision

A private report should help executives decide what to improve, test, stop or verify. It should state scope, show disagreements and propose a bounded action with an owner and measure. It must remain useful even if research does not support a single headline score.

Development requires content review, bilingual cognitive work, feasibility testing, independent coding and empirical evaluation for declared

purposes. Sample requirements follow the construct, comparisons and precision, not a convenient target. A 40-60-company feasibility plan is a capacity assumption, not validation power. [Best Practices for Developing and Validating Scales for Health, Social, and Behavioral Research: A Primer](#).

Public research describes source-backed examples and limits of a collection. A future participant study would examine permissioned evidence over a defined period. Their scales cannot be converted, and the public series cannot certify enterprise capability.

Leadership agenda

Executive sponsor: identify the recurring decision the assessment should improve.

Method lead: specify the construct and evidence needed for that use.

Operating team: retain outcomes, interventions and uncertainty alongside successes.

Scope and reading boundaries

The source corpus contains 80 purposefully selected entity records across eight sectors, ten per sector. Headquarters markets are Saudi Arabia, the UAE, Qatar, Kuwait and Bahrain; Oman is outside this edition. Groups, subsidiaries and specific operations have different scopes. This is not a representative sample. Source cutoff: 10 October 2026; publication edition: 11 October 2026.

Source extraction and challenge were AI-assisted, not independent human assurance or verification of current operations. Source visibility is not adoption prevalence. Unreported details are unresolved questions, not absent capabilities. Named examples imply neither participation, endorsement, confidential access nor a Mayden client relationship.

EXHIBIT 02 / CONCEPTUAL TOOL

Keep capability, outcomes and evidence on separate ledgers

Mayden analytical framework / questions to test, not measured maturity stages

Practices	Record the work design, controls, human roles and recurring use.
Outcomes	Record quality, cost, risk and service consequences against an appropriate comparator.
Evidence	Record provenance, access, dates, limitations and the scope of assurance.


Mayden conceptual framework. These questions specify evidence to request; they do not assert that a named organisation has or lacks the practice. This is not a validated instrument or an observed causal sequence.

Unit: Conceptual decision or evidence request. Research cutoff: 2026-10-10.

EXHIBIT 03 / CONCEPTUAL TOOL

A public source and a private assessment answer different questions

Mayden analytical framework / questions to test, not measured maturity stages



Mayden conceptual framework. These questions specify evidence to request; they do not assert that a named organisation has or lacks the practice. This is not a validated instrument or an observed causal sequence.

Unit: Conceptual decision or evidence request. Research cutoff: 2026-10-10.

05 / THE LEADERSHIP AGENDA

Ask for evidence that changes the decision.

01 / Executive sponsor

Evidence to request: identify the recurring decision the assessment should improve.

02 / Method lead

Evidence to request: specify the construct and evidence needed for that use.

03 / Operating team

Evidence to request: retain outcomes, interventions and uncertainty alongside successes.

Method and boundaries

This is an AI-assisted, bounded public-source review of a purposefully selected corpus. It publishes sector research and attributed operating examples, without current individual-company scores, ranks or coded profiles. Source accounts establish what a publisher reports; they do not independently assure operations. The working definition and nine candidate themes require empirical validation.

Recorded source dates run from 2019-05-14 to 2026-10-01. 79 of 224 qualifying claim records have no recorded publication date. This is not a common observation window or confirmation of continuing use.

The confidential opt-in diagnostic remains in development. Independent human review, instrument testing, participant data collection and comparable benchmark norms are future gates. Prior named public-evidence comparisons remain explicitly historical; changes in publication format do not establish changes inside companies.

[Inspect the method](#) / [Explore sector research](#)

06 / SOURCES & PROVENANCE

Follow the argument back to its source.

The following sources support the cited analysis or attributed examples. Inclusion is not independent verification, company endorsement or assurance. Research cutoff: 10 October 2026. Publication revision: 11 October 2026.

S01 / Generative AI at Work

Quarterly Journal of Economics / 2025-02-04 / peer-reviewed empirical

Locator: Staggered workplace rollout; observational causal design

[Open source](#)

S02 / Experimental evidence on the productivity effects of generative artificial intelligence

Science / 2023-07-13 / peer-reviewed experiment

Locator: Preregistered randomized online experiment

[Open source](#)

S03 / Navigating the Jagged Technological Frontier

Harvard Business School working paper 24-013 / 2023-09 / working-paper experiment

Locator: Randomized field experiment

[Open source](#)

S04 / Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity

METR / 2025-07-10 / original experiment

Locator: Task-randomized controlled trial

[Open source](#)

S05 / We are Changing our Developer Productivity Experiment Design

METR / 2026-02-24 / original follow-up methods report

Locator: Follow-up randomized study; methodological disclosure

[Open source](#)

S06 / Conventional Wisdom on Measurement: A Structural Equation Perspective

Psychological Bulletin / 1991-09 / peer-reviewed-methodological

Locator: Conceptual measurement-model analysis

[Open source](#)

S07 / Handbook on Constructing Composite Indicators: Methodology and User Guide

OECD / European Union / EC-JRC / 2008-08-22 / measurement guidance

Locator: Methodological handbook

[Open source](#)

S08 / Best Practices for Developing and Validating Scales for Health, Social, and Behavioral Research: A Primer

Frontiers in Public Health / 2018-06-11 / peer-reviewed measurement guidance

Locator: Methodological primer

[Open source](#)

About Mayden

Mayden connects senior judgement, research and engineering to help organisations turn strategic choices into working capabilities. The AI Native Index is a research programme on recurring work, decisions and accountability.

[Explore the Index](#) / [Submit a correction](#)