




MANUFACTURING & MATERIALS

From process signals to operating decisions

Production, maintenance and quality across ten selected industrial entities.



ORIGINAL MAYDEN GRAPHIC / EVIDENCE IN MOTION

01 / THE EXECUTIVE VIEW

Start with the operating evidence.

Production, maintenance and quality across ten selected industrial entities.

10

selected entity records

17

recorded claims

17

eligible claims

5

unranked records

01 / WHAT THE RECORD SHOWS

Start at the production decision. EGA reports AI vision for carbon-anode inspection at Al Taweelah. Ask how detected defects affect operator action, rework and accepted quality.

Claims: [MAN-EGA01](#)

02 / WHAT THE RECORD SHOWS

Separate a completed model from sustained use. Alba describes a density-prediction model; EGA describes a planned autonomous-potline pilot. Confirm production use and decision authority before interpreting either as autonomous manufacturing.

Claims: [MAN-ALBA01](#) / [MAN-EGA02](#)

03 / WHAT THE RECORD SHOWS

Measure the planning workflow end to end. Ducab reports AI within manufacturing order planning. Examine accepted schedules, planner overrides and actual operating results; improvement targets are not achieved outcomes.

Claims: [MAN-DUCAB01](#)

Public-evidence review. Selected entity scopes. No company participation, population inference or independent assurance implied.

Read the argument, inspect the exhibits, then use the leadership questions to identify the next evidence your organisation should produce. All report bodies are in English; Arabic executive summaries are available online.

EXHIBIT 01 / SCOPED CASE COMPARISON

Three factory decisions. Three evidence tests.

Inspection, safety prioritisation and planning change different operating objects.

Inspect quality / Carbon-anode quality decision	
Source account. EGA reports real-time AI vision checking carbon-anode quality at AI Taweelah. Boundary. Site-specific inspection; group-wide autonomy not established.	Mayden question. What is rejected, reviewed or reworked after the signal?
Evidence: MAN-EGA01	
Prioritise safety attention / High-risk operating area	
Source account. Alba describes AI using incident patterns and risk drivers to identify higher-risk areas. Boundary. Identification is reported; causal safety improvement is not verified.	Mayden question. What action changes, and how are false alerts and missed risks reviewed?
Evidence: MAN-ALBA04	
Plan production / Manufacturing order schedule	
Source account. Ducab states that AI software is used in order planning within Blade. Boundary. Facility coverage undisclosed; targets are not achieved outcomes.	Mayden question. Which schedule changes are accepted, overridden and measured?
Evidence: MAN-DUCAB01	

Source: selected claims in PE-2026.10-v4. Company and supplier reports are attributed accounts. The questions are Mayden editorial interpretations, not measured sector results or independent assurance.

EXHIBIT 02 / PUBLIC EVIDENCE

Inspect the operating profile, not only the total

Highest eligible anchor per theme; U = no qualifying evidence located

	Core	People	Partners	Interface	Lineage	Testing	Authority	Knowledge	Control	Points
Aluminium Bahrain (Alba)	3	3	1	3	U	1	U	1	2	52
SABIC	3	2	U	U	U	1	U	2	2	37
Emirates Global Aluminium	3	U	U	U	U	1	U	U	U	15
Maaden	2	U	2	U	U	U	U	U	U	15
Ducab	2	U	U	U	U	U	U	U	U	7
EQUATE	U	U	U	U	U	U	U	U	U	U
GARMCO	U	U	U	U	U	U	U	U	U	U
Industries Qatar	U	U	U	U	U	U	U	U	U	U
Petrochemical Industries Company	U	U	U	U	U	U	U	U	U	U
QAPCO	U	U	U	U	U	U	U	U	U	U


Unknown
Commitment / pilot
Implemented
Operating detail

Source: Mayden PE-2026.10-v4. Each row is a named entity scope. U does not mean absent capability. Anchors are ordinal evidence-specificity judgements; scopes and disclosure access differ.

EXHIBIT 03 / PUBLIC EVIDENCE

Documented evidence points in the selected cohort

Fixed denominator of 27; normalized to 0-100; ties use unrounded totals



Source: frozen release calculation. These are public-evidence points, not a maturity percentage, operational audit or population ranking. No-evidence records stay unranked.

02 / ANALYSIS

Manufacturing & Materials – Mayden AI Native Index research edition

Evidence cutoff: 10 October 2026. Ten purposefully selected entities across five markets. Public evidence, not a measurement of actual enterprise maturity.

A factory can announce an AI partnership without changing a production decision. It can also run a useful predictive model for years without describing it publicly. An index that treats either possibility carelessly will reward communications activity and punish silence. The manufacturing research therefore asks a narrower, more useful question: what operating practice can an external reader actually trace to a primary source, and how far does that source support the claim?

The strongest cases in this corpus describe a decision close to production: identifying an anode-quality issue, prioritising a safety risk, or using AI in an order-planning workflow. These are more informative than an undifferentiated count of digital initiatives. They reveal where information enters work. They do not establish that the whole enterprise is AI-native, that a model is accurate, or that a reported improvement was caused by AI.

This report turns those cases into questions executives can use. Its central proposition is an editorial interpretation: an AI-native manufacturer would need repeatable arrangements for deploying, supervising, maintaining and learning from models across operating conditions. A successful isolated application is evidence of an application. The additional organisational arrangements remain an empirical question.

What the cohort covers

The ten entities are SABIC and Maaden in Saudi Arabia; Emirates Global Aluminium and Ducab in the UAE; Industries Qatar and QAPCO in Qatar; EQUATE and Petrochemical Industries Company in Kuwait; and Aluminium Bahrain and GARMCO in Bahrain.

Selection creates a small, purposeful view of chemicals, mining, metals and industrial manufacturing. It is not a census, a sample of the largest ten firms, or a representative estimate of GCC adoption. Oman is outside this five-market expansion.

Corporate boundaries are consequential. Industries Qatar has exposure to QAPCO; PIC has an ownership relationship with EQUATE. These are not independent observations of entirely separate economic systems. Maaden's investment in Alba also creates overlap, while SABIC has an ownership relationship with an energy-sector company covered elsewhere in the Index. No operating claim is transferred from an investee to its shareholder simply because that relationship exists. Alba's retained cases concern its Bahrain operation or a specified central function, rather than every operation in its changing corporate perimeter.

The cohort contains both integrated operators and investment-group structures. Comparing their disclosure evidence can be useful if that distinction is visible. Comparing the resulting numbers as if they represented identical factories would be misleading. Readers should open the company scope and supporting passage before interpreting a difference.

Production evidence is strongest when the unit is visible

SABIC's 2025 report states that its AI and predictive tools are deployed across a majority of global manufacturing plants. It separately describes pilots in demand forecasting, pricing and planning; connected laboratory systems; workforce adoption programmes; and risk assessments that include AI. The separation matters: a broadly deployed production-tool portfolio does not turn each commercial pilot into a production system. Nor does it establish autonomous control. [SABIC Integrated Annual Report 2025, digital and cybersecurity sections, printed pages 60–62](#)

The executive lesson is to insist on a denominator with a stable definition. "Most plants" is more informative than "many use cases," but the next discussion still needs to identify eligible sites, processes covered, usage frequency and retirement of unsuccessful models. A portfolio count can rise because a company creates more prototypes, because it divides applications differently, or because it has genuinely expanded productive use. Only the last interpretation necessarily changes operating capability.

EGA provides a more localised example. Its January 2025 lighthouse announcement describes implemented AI vision for real-time carbon-anode quality assessment. That is a named production task with a recognisable observation mechanism. The surrounding announcement also discusses a broader Industry 4.0 portfolio. Those portfolio totals and general digital-skilling numbers should not be relabelled as AI totals. [EGA announcement, 15 January 2025](#)

A quality model creates several different questions. Does it flag an item for inspection, change a process setting, or reject output automatically? Who resolves disagreement between the model and an experienced operator? What happens when raw-material characteristics change? The public case supports the existence of a particular application; these questions define the deeper evidence needed for an organisational benchmark.

Ducab's Blade announcement identifies AI software used in manufacturing-order planning. The same release describes expected improvements in production and downtime. This review retains the stated planning use and keeps those expected benefits separate. It does not present targets as achieved results or treat every sensor in the wider project as AI. [Ducab Blade announcement, 13 June 2023](#)

For a planning application, usefulness depends on more than forecast accuracy. An apparently better plan may violate an equipment constraint, create downstream inventory, or increase schedule instability. A serious company interview should connect the recommendation to the supervisor who

accepts it and the production consequences that person monitors.

A pilot deserves recognition without being called a rollout

EGA's December 2024 EX announcement is explicit about an experimental facility and an intended future autonomous production system. It describes the construction of pilot reduction cells, prospective AI capabilities and a longer industrialisation horizon. Existing developed use cases are discussed alongside these ambitions. The future autonomous-control objective remains a different claim from an implemented quality-inspection application. [EGA EX pilot announcement, 9 December 2024](#)

This distinction is useful even when a later announcement confirms that a physical pilot has started producing metal. Commissioning the industrial asset does not, by itself, establish that a particular AI control is operating routinely. The evidence needed to cross that boundary is a dated account of the AI system entering service, its operating envelope and the accountable owner.

Alba's 2023 announcement describes completion of a green-anode density prediction model and a dashboard developed with NAIRDC. The review treats this as a specific completed development or pilot rather than assuming sustained production deployment. Separately, the 2025 Seeq announcement contains future integration language despite its adoption-oriented headline. It is retained as a named commitment. [Alba-NAIRDC project, 12 July 2023](#); [Alba-Seeq announcement, 27 April 2025](#)

Executives should not read a lower evidence stage as criticism of experimentation. A well-run pilot can prevent an expensive operational mistake. The relevant management question is whether there is a disciplined transition: a baseline, a success criterion, a safety review, an owner, a decision date and a route to withdrawal. A benchmark should be able to recognise that discipline without pretending that intent has already produced enterprise-wide change.

The organisation appears in the less theatrical applications

Alba's first-quarter 2026 company magazine describes AI used to identify high-risk operating areas from incident patterns and risk drivers. It also describes a Scope of Work assistant that helps prepare structured contract documents for an approval process. A conduct-policy chatbot is described with less conclusive deployment language. These are different levels of support, even though they appear in the same publication. The safety narrative does not establish a causal reduction in accidents; the drafting assistant does not establish autonomous contracting. [Alba AI Bayan Q1 2026, physical pages 24-25](#)

These examples direct attention beyond the furnace or production line. An enterprise changes when a tool becomes part of how work is specified, reviewed and escalated. Contract drafting can affect procurement quality. Safety prioritisation can influence the allocation of supervisory attention. Both deserve scrutiny because their consequences depend on human interpretation and follow-through.

The next interview should ask for a recent decision reconstructed end to end. Which information entered the tool? Which recommendation emerged? Who inspected it? What changed as a result? Could the organisation reproduce the reasoning after a dispute or incident? A useful answer would identify records and responsibilities. A generic assurance that humans remain in control would leave the most important details unresolved.

Alba's August 2024 course announcement provides another bounded organisational signal: a three-day AI programme underway for 35 employees, with specified generative-AI and prompting topics. Those are participants in an announced course, not independently verified completions or evidence that the workforce has mastered AI. [Alba training announcement, 21 August 2024](#)

A training measure should distinguish attendance, assessed competence and changed work. Different roles need different capabilities. A process engineer

may need to recognise model drift; a procurement specialist may need to verify generated specifications; a manager may need to judge when automation is inappropriate. Counting all three people as identically trained loses information that matters to performance.

Alba also disclosed an updated Code of Conduct with responsible-AI guidance in September 2026. This supports the presence of an issued policy. The disclosure does not supply enough detail to assess whether particular model controls have been applied effectively. [Alba regulatory disclosure, 23 September 2026](#)

Exploration is a distinct operating setting

Maaden's technology and partnership pages describe AI and data modelling in mineral exploration, including work with Ivanhoe Electric. They are useful primary descriptions of an operating activity and a named collaboration. They are live undated pages, however, and do not provide a deployment date, model-performance evaluation or a clear coverage denominator. [Maaden technology page; Maaden partnerships](#)

Exploration is not the same decision environment as continuous process control. Its evidence can involve sparse observations, uncertain geological interpretation and long feedback cycles. A common organisational framework needs sector-specific questions. The principle is to compare how an organisation makes and governs AI-supported decisions, while resisting the temptation to compare unlike model metrics.

An exploration team could explain how recommendations change survey priorities, how new field observations are incorporated, and how unsuccessful predictions are documented. A smelter could explain how a model reacts to process changes and how operators override recommendations. These are different technical accounts of a common organisational obligation: keeping a decision system reliable in use.

What remains unknown is part of the result

Industries Qatar's examined annual-report passages mention AI in the economic and market backdrop. Such language does not demonstrate an implemented group practice. QAPCO's report archive and older PDF links presented access problems; indexed references to a predictive-maintenance plan were not promoted into verified source evidence. EQUATE's accessible executive summary did not yield a qualifying AI practice in the passages searched. For PIC and GARMCO, the bounded company-site search and identity material did not provide qualifying operating evidence.

These findings do not establish that those companies lack AI. They establish the limits of this research corpus. Report access failures, different disclosure customs, corporate structure and search vocabulary can all affect visibility. The correct label is unknown within the bounded search, accompanied by the search log.

A reader must also distinguish a failed retrieval from a source that was read and found inapplicable. The former leaves an unresolved lead. The latter may rule out a particular interpretation, such as treating a general technology investment as AI. The dataset records these differences so that a subsequent researcher can make a targeted correction rather than repeat the entire search.

Questions that can turn disclosure research into a benchmark

A proprietary company assessment should start with two material decisions selected by the firm and one selected by the interviewer. This reduces the chance that the interview covers only a flagship demonstration. For each decision, the respondent should identify the task, operating owner, affected sites, user population, deployment date and current stage. The assessment should then request a small evidence sample rather than unrestricted access to operational systems.

The questions should cover nine dimensions with sector-specific anchors. Where is AI materially involved in core work? Which employees have demonstrated role-relevant competence? What responsibilities sit with technology partners? How do outputs enter systems or work queues? Can a material decision be reconstructed from retained records? What evidence determines whether an experiment proceeds? Which actions can the system take, and who can stop it? How is operational knowledge maintained? Which AI-specific controls have actually been applied?

These are proposed interview questions, not findings about the cohort. They should be tested with operators for clarity before being used to create a validated organisational measure. A polished questionnaire alone cannot establish validity. Mayden would still need consistent coding, independent adjudication, evidence of reliable responses and investigation of whether the resulting dimensions relate to meaningful external outcomes.

Reading the manufacturing evidence exhibits

The heatmap compares the ten selected entities across nine dimensions. U means no qualifying evidence was located in the bounded corpus; it does not mean the capability is absent. Level 1 denotes a specific commitment or pilot, level 2 an implemented practice, and level 3 an implemented practice with concrete operating detail. Each cell takes the highest supported, adjudicated level for that dimension. Several descriptions of one application do not accumulate extra points.

The points comparison uses a fixed denominator of 27: nine dimensions with a maximum of three points each. The displayed value is the supported total divided by 27 and multiplied by 100, rounded to a whole number. Unknown cells add no documented points, but remain unknown in interpretation. A higher total therefore describes a stronger public evidence record under these rules; it is not a percentage of enterprise maturity or proof of better factory performance. Equal underlying totals share a position, and records without qualifying evidence remain unranked.

The anchor-count exhibit separates commitments or pilots, implemented practices, and implemented practices with operating detail. These are claim counts, not counts of independent projects, factories or companies. One operating application may support more than one dimension. This helps explain the distinction between an implemented inspection practice and a future autonomous-control objective without treating every mention of AI as equivalent.

The source-dependence exhibit removes one source URL at a time and recalculates the points, showing the result when the most influential URL is removed. A large decrease identifies concentration in the disclosed evidence, often in an annual report. It does not simulate a factory losing its AI capability. Different URLs can also contain related company assertions, so a smaller decrease is not proof of independent corroboration. Read the chart alongside the entity scopes: shareholder and investee records can overlap even when their source URLs differ.

The practical outcome of this edition is an inspectable research base and a sharper conversation with manufacturers. A future

benchmark becomes credible by preserving these distinctions over time: company versus site, intention versus operation, activity versus outcome, and undisclosed versus absent. The index should make it easier to ask the next difficult question, not make the available evidence appear more complete than it is.

Evidence and reproducibility. The companion manufacturing-evidence.json contains company scopes, source URLs, passage locators, proposed evidence levels and per-company search limits. Separate AI-assisted source challenge is recorded separately; the released calculations use the retained adjudicated decisions. Company assertions remain company assertions. This edition includes no confidential company access, respondent survey, statistical meta-analysis or validation of actual AI-native organisational maturity.

Review transparency. Extraction and source challenge in this edition were AI-assisted. Separate AI readers do not constitute human expert sign-off or independent assurance. Company claims, research findings and Mayden interpretations retain their distinct scope.

EXHIBIT 04 / INSPECT THE ASSUMPTION

Separate a commitment from a deployed practice

Eligible claim records, n = 17; one practice may support multiple distinct themes



Source: claim ledger. 0 excluded records are outside these bars. Counts are claims, not independent systems, companies or estimates of adoption.

EXHIBIT 05 / INSPECT THE ASSUMPTION

How much of the profile rests on one retained source URL?

Baseline points versus points after removing the most influential primary source URL



Source: deterministic URL-removal check. Remaining URLs can repeat the same disclosure; required corroborating-source dependencies are not removed. This is not a confidence interval or a test of source truth.

03 / A CLOSER LOOK

Aluminium Bahrain (Alba)

Aluminium Bahrain B.S.C.; retained evidence concerns Bahrain smelter and named central functions, not all newly acquired operations. Maaden shareholding creates overlap.

ELIGIBLE ANCHOR 3 / MAN-ALBA02

Alba reports a three-day AI course underway for 35 employees, covering generative AI and prompting.

Boundary. Claim scope: Aluminium Bahrain B.S.C.; retained evidence concerns Bahrain smelter and named central functions, not all newly acquired operations. Maaden shareholding creates overlap.. 35 participants are not 35 certified completions; no productivity effect established. Application: internal-operating-practice.

Locator. Body paragraphs on participants, duration and course topics

[Inspect the source](#)

ELIGIBLE ANCHOR 3 / MAN-ALBA04

Alba describes deployed AI using incident patterns and risk drivers to identify high-risk operating areas.

Boundary. Claim scope: Aluminium Bahrain B.S.C.; retained evidence concerns Bahrain smelter and named central functions, not all newly acquired operations. Maaden shareholding creates overlap.. No verified predictive accuracy or causal safety improvement. Application: internal-operating-practice.

Locator. Physical24, printed46-47, AI safety article

[Inspect the source](#)

The question for leadership

Locate the production constraint.

Company/group scope follows each cited claim. Headquarters is not the geography of all operations. Public disclosures and search completeness differ; no national or population inference.

03 / A CLOSER LOOK

SABIC

SABIC group global manufacturing; claim scope restricted to named functions. Aramco ownership creates cross-sector overlap.

ELIGIBLE ANCHOR 3 / MAN-SABIC01

AI process and predictive analytics tools are deployed at a majority of global manufacturing plants.

Boundary. Claim scope: SABIC group global manufacturing; claim scope restricted to named functions. Aramco ownership creates cross-sector overlap.. Mixed predictive analytics/AI tool portfolio; not every tool is AI or every plant covered. Application: internal-operating-practice.

Locator. Digital Transformation, printed60 / physical60, deployment paragraph

[Inspect the source](#)

ELIGIBLE ANCHOR 2 / MAN-SABIC03

GenAI platforms and AI infrastructure connected previously isolated laboratory systems.

Boundary. Claim scope: SABIC group global manufacturing; claim scope restricted to named functions. Aramco ownership creates cross-sector overlap.. No documented knowledge maintenance, retrieval evaluation or lineage controls. Application: internal-operating-practice.

Locator. Digital Transformation, printed60 / physical60, R&D paragraph

[Inspect the source](#)

The question for leadership

Locate the production constraint.

Company/group scope follows each cited claim. Headquarters is not the geography of all operations. Public disclosures and search completeness differ; no national or population inference.

03 / A CLOSER LOOK

EQUATE

EQUATE group; not unrelated academic projects named EQUATE. PIC equity ownership creates overlap.

No claim qualified in this bounded review. An empty public record is a research gap, not evidence that AI capability is absent.

The question for leadership

Locate the production constraint.

Company/group scope follows each cited claim. Headquarters is not the geography of all operations. Public disclosures and search completeness differ; no national or population inference.

04 / THE LEADERSHIP AGENDA

Ask for evidence that changes the decision.

01 / Locate the production constraint.

Evidence to request: Named line/site, quality or maintenance decision and accepted operating baseline.

02 / Inspect integration with plant control.

Evidence to request: Action authority, operator override, validation and safety boundaries.

03 / Distinguish a useful model from a scaled practice.

Evidence to request: Utilization, covered assets, retraining ownership and reproducible results.

Method and boundaries

This is an AI-assisted, bounded public-source review using a purposeful corpus. Nine themes retain their highest eligible anchor: 1 for a specific commitment or pilot, 2 for an implemented practice, 3 for implemented practice with operating detail. U means unknown. Points equal 100 times the sum of anchors divided by 27. Ordinal steps and equal weights are design assumptions, not validated importance estimates.

Reported outcomes, actual company capability and disclosure depth are different objects. Source-linked corporate or vendor statements are not independently assured operations. Deterministic sensitivity tests are not confidence intervals. Sector scopes, reporting periods and access differ; no national rankings, population prevalence or causal inference are made. The confidential company instrument remains a separate development track.

Dated sources run from 2023-06-13 to 2026-09-23. 2 of 17 eligible claim records have undated sources. Publication dates do not confirm continuing operation or a common observation period; undated does not mean stale.

[Inspect the full method](#) / [Reproduce the evidence](#)

05 / SOURCES & PROVENANCE

Follow the argument back to its source.

Sources below include eligible and excluded research leads. Inclusion in this list does not make a source eligible or independently verified. Exact claim decisions, source locators and limitations are retained in the online evidence ledger. Research cutoff: 10 October 2026.

S01 / Alba and NAIRDC artificial intelligence collaboration

Alba / 2023-07-12 / corporate-operating-announcement

Locator: Body paragraphs on completed Green Anode Density Prediction project

[Open source](#)

S02 / Alba partners with NVTC R&D arm to upskill employees

Alba / 2024-08-21 / corporate-operating-announcement

Locator: Body paragraphs on participants, duration and course topics

[Open source](#)

S03 / Alba adopts AI-powered Seeq platform

Alba / 2025-04-27 / corporate-operating-announcement

Locator: Body paragraphs on agreement and future integration

[Open source](#)

S04 / AI Bayan Q1 2026

Alba / 2026 / company-magazine

Locator: Physical24, printed46-47, AI safety article

[Open source](#)

S05 / Updated Code of Conduct disclosure

Alba / 2026-09-23 / regulatory-disclosure

Locator: Physical2, launch and new guidance paragraphs

[Open source](#)

S06 / SABIC Integrated Annual Report 2025

SABIC / 2026 / annual-report

Locator: Digital Transformation, printed60 / physical60, deployment paragraph

[Open source](#)

S07 / EGA designated Industry 4.0 Global Lighthouse

EGA / 2025-01-15 / corporate-operating-announcement

Locator: Body paragraph describing real-time carbon-anode quality assessment

[Open source](#)

S08 / EX technology pilot and Industry 4.0 capabilities

EGA / 2024-12-09 / corporate-operating-announcement

Locator: Body paragraphs on ten pilot reduction cells and autonomous-control goal

[Open source](#)

S09 / Tech innovations

Maaden / Undated / corporate-operating-page

Locator: Exploration technology paragraph

[Open source](#)

S10 / Partnerships

Maaden / Undated / corporate-operating-page

Locator: Ivanhoe Electric partnership paragraph

[Open source](#)

S11 / Ducab launches Blade smart factory project

Ducab / 2023-06-13 / corporate-operating-announcement

Locator: Body paragraph on AI software and planning orders

[Open source](#)

About Mayden

Mayden connects senior judgement, research and engineering to help organisations move from strategic choices to working capabilities. This research programme examines how AI changes recurring work, decisions and accountability.

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