



Japan publishes, and gets cited.

83 companies here reach the
published APAC top 1,000.



Authority, freshly earned

Of the five signals the index reads in Japan, freshness and authority is the one its record leans on hardest, and it leans further on it than on any of its other four: what its companies publish is current, and it gets cited. 232 leave enough of a public trail to score and 83 reach the published APAC top 1,000. Manufacturing carries the visible end with 47 ranked companies. Proof of capability is the quietest signal here, and the 17 companies publishing any average 62 of 100, so what is missing is the number publishing rather than the strength of what they show.

Every figure here measures what a company shows, not what it can do. That distance is the opportunity.

What's inside

Executive summary	3
Why visibility is the asset	4
Who is publishing proof	5
How the evidence grades	6
Where the signal is strong, and where it is not	7
Signal fingerprint	8
Readiness, and who is visible	9
Four tiers of visibility	10
The visibility ladder	11
Where the sectors stand	12
The national champions	13
How these numbers were built	14
What this report does not claim	15
What to do with a visibility score	16
About AIR APAC	17
Sources and notes	18
Where the build differs from the method	19



58

Readiness score for Japan. Range 52-64.



44.2

Average visibility score across 232 scored companies. Public evidence only.



83

Companies here that reach the published APAC top 1,000, of 232 scored.

Source: AIR APAC AI Readiness Pulse, Q2 2026 (v0.2A). Scores measure public evidence, not verified capability. A company score says what it publishes, never what it can do.

EXECUTIVE SUMMARY

Current, and widely cited

Readiness 58 of 100, range 52-64.

Visibility 44.2 of 100 across 232 scored companies.

Japan at a glance, Q2 2026

58

Readiness score out of 100.
Range 52-64, evidence graded B.

29.8%

Of 346 companies screened in Japan,
too little public evidence to rank.

A low score here reflects the public record, not a company's capability. Japan is among the record leaders: of the five signals the index reads, its record leans hardest on freshness and authority and thinnest on capability visibility. The conditions companies here operate under score 58, and what those companies publish scores 44.2. Of the two, the second is the one companies here control, and the quicker to change: 93% of the 232 companies with enough of a public trail to score publish no proof of capability at all.

Japan's record is current and it is cited. The 17 companies that also publish proof of capability average 62 of 100, on a method the rest can copy.

Exhibit 1. Tier mix in Japan

Emerging Findable Building Proven

Japan

64%

20%

12%

Source: AIR APAC AI Readiness Pulse, Q2 2026 (v0.2A). Scores measure public evidence, not verified capability.



83 companies in Japan reach the published APAC top 1,000.
149 of 232 are scored already, next in line for the published list.

Invisible is not the same as absent

Buyers, investors and partners treat it that way. Most decisions about a company are shaped before anyone is contacted, from what they can find. The record is the only version of it most people ever meet. and here most of the market already leaves one, which moves the question from whether a record exists to what it says about capability. In Japan, 103 of 346 companies screened leave too little behind to assess, which is the addressable group rather than the verdict.



Capital and customers screen before they call

Investors, corporate buyers and procurement teams build a shortlist from public material long before a meeting exists. A company absent from that material is not rejected on the merits. It never reaches the shortlist.



Partners and talent read the same record

Prospective partners, joint-venture counterparties and senior candidates all check what an organisation says about its own AI work. Silence is read as inactivity, which is often wrong and rarely without cost.



AI assistants now answer the question first

When someone asks an AI system what a company is doing with AI, the answer is assembled from public sources. A company that publishes nothing cedes its own description to whatever else the model finds.

"Most companies we look at are doing more with AI than their customers, investors and partners can see. The work is real, the evidence of it is thin, and people decide on the evidence."

Indhran Seghar
Founding Director, AIR APAC



83

Companies here reach the published APAC top 1,000.



71%

Of 232 scored rest on grade A or B evidence.



149


Scored already, and next in line for the published list.

Who is publishing proof

17 companies in Japan publish capability evidence; the rest do not yet.

Exhibit B. Proof concentration in Japan.

17 publish proof, avg 62/100



Source: AIR APAC AI Readiness Pulse, Q2 2026 (v0.2A). Scores measure public evidence, not verified capability.



Proof changes the quality of the record.

Capability evidence shows what was built or deployed. It gives buyers, investors and partners something concrete to evaluate, rather than another statement of intent.



7%

Of ranked companies publish capability evidence.



62

Average visibility score among those publishing proof.




215

Ranked companies have not published capability proof yet.

How the evidence grades

How the evidence behind Japan's ranked companies grades out.

Exhibit C. Evidence quality mix in Japan.



Source: AIR APAC AI Readiness Pulse, Q2 2026 (v0.2A). Scores measure public evidence, not verified capability.



A grade describes traceability, not capability.

Higher grades mean the public claim is easier to verify through first-party or attributable evidence. They do not certify that the underlying AI system performs well.



71%

Of ranked companies rest on grade A or B evidence.



16

Ranked companies carry the strongest evidence grade.




29%

Of ranked companies rely on grade C or D evidence.

Japan's strongest and weakest public signals

The five signals are not on one scale. Nearly every company in the index is discoverable and nearly none publishes proof of what it can do, so the longest bar below is a property of the index before it is a property of Japan. Read against its own other four signals, this record leans hardest on freshness and authority and thinnest on capability visibility. The first is where companies here have already shown they can publish; the second is where the same effort has not yet gone, and it is the clearest place to start.

Exhibit 2. Average score by signal in Japan



Source: AIR APAC AI Readiness Pulse, Q2 2026 (v0.2A). Scores measure public evidence, not verified capability.

WHAT THIS MEANS

The lean is freshness and authority.


17 companies in Japan publish capability evidence, averaging 62 and well clear of the market. It works where it exists. The other 93% have yet to publish it, against 83% regionally.

Page 17 sets out what closing it takes: dated case studies, named leads, AI roles on a careers page. Averages cover ranked companies only. Work published only in Japan's own language is understated here, not absent.

Signal fingerprint

Japan's signal profile, dimension by dimension.

Exhibit D. Signal shape, Japan's five dimensions.



Source: AIR APAC AI Readiness Pulse, Q2 2026 (v0.2A). Scores measure public evidence, not verified capability.

Japan's readiness, and who is visible

Japan scores 58, within a range of 52-64 on evidence we grade B. Read it as this market's own country context, not as a position against its neighbours: markets were not sourced to a common depth, so a gap between two of these scores does not have a reliable direction. The range is set by how much country evidence we can verify, and it narrows as that coverage improves.

Exhibit 3. Japan's readiness score and range, out of 100



Source: AIR APAC AI Readiness Pulse, Q2 2026 (v0.2A). Scores measure public evidence, not verified capability.

The ten most visible companies in Japan

APAC	Company	Score	Grade	Sector
#22	LegalOn Technologies	71	B	Tech/Fintech
#24	Renesas Electronics Corporation	71	B	Manufacturing
#53	MUFG Bank, Ltd.	68	A	Financial Services
#68	ITOCHU Techno-Solutions Corporation	66	B	Professional Services
#72	Mercari	66	B	Retail/E-commerce
#73	Mitsubishi Electric Corporation	66	B	Manufacturing
#74	NS Solutions Corporation	66	B	Professional Services
#75	Preferred Networks, Inc.	66	B	Tech/Fintech
#78	SCSK Corporation	66	B	Professional Services
#99	Mitsubishi Research Institute	65	A	Professional Services

Ranks are region-wide, not Japan-only. Score is the visibility score (RPS in the dataset). Evidence grades run A to D. Source: AIR APAC AI Readiness Pulse, Q2 2026 (v0.2A). Scores measure public evidence, not verified capability.

The most visible company here ranks 22nd in the region. Professional Services companies take 4 of the 10 highest individual places, more than any other sector. By sector average, though, manufacturing leads at 47.7.

Four tiers, and where companies in Japan sit

A tier is a position in the published regional ranking, not a grade anyone awards. It says how much of a company's AI work a buyer or an investor can verify, and nothing at all about whether that work is any good.



Tier 1
Proven

10

Region rank 1-100. Publishing enough, recently enough, and across enough of the five signals that the public record alone establishes an AI position.



Tier 2
Building

27

Region rank 101-400. A consistent public trail, usually strong on two or three signals and thin on evidence of hiring.



Tier 3
Findable

46

Region rank 401-1,000. Credibly active and possible to verify, but sparse enough that it takes a deliberate search to find.



Tier 4
Emerging

149

Ranked beyond 1,000. Clears the evidence threshold but does not reach the published list. Real companies, scored and unlisted, and the largest group in almost every market.

10 of the 232 companies in Japan that clear the threshold reach tier 1, and 149 are emerging, 64% of those scored. Emerging is the normal state of the region, and the group standing closest to the tier above.



Moving up a tier is a disclosure decision, not a technology one, because every signal this index reads is something a company already controls.



71%

Of scored companies rest on grade A or B evidence.



83

Reach the published regional top 1,000, of 232 scored.



29.8%

Of 346 screened leave too little to score.

The visibility ladder

From 346 screened to 83 published: the funnel behind Japan's ranked table.

Exhibit E. The visibility ladder in Japan.



Source: AIR APAC AI Readiness Pulse, Q2 2026 (v0.2A). Scores measure public evidence, not verified capability.



Each step is a different filter, not a ranking.

Screened is the sourced universe. Scored is what left enough public trail to read. Listed merges a group scored under several web domains into one row. Published is the regional top 1,000.



346

Companies sourced in this market before any filter applies.



11

Rows merged where one group was scored under several web domains.



36%

Of listed companies reach the published regional top 1,000.

Where the sectors stand

The five sectors with the strongest average visibility in Japan.

Exhibit F. Sector strengths in Japan.



Source: AIR APAC AI Readiness Pulse, Q2 2026 (v0.2A). Scores measure public evidence, not verified capability.



A sector score is a publishing habit, not an adoption rate.

It says how visibly an industry writes up its AI work in public, English-language sources.

It does not say how much of that work is under way.



47.7

Mean visibility in Manufacturing, the strongest of these five.



5.2

Points separate that sector from the fifth on this page.




204

Ranked companies sit inside these five sectors.

The national champions

The ten most visible companies in Japan, coloured by evidence grade.

Exhibit G. National champions in Japan, top ten by visibility.



Source: AIR APAC AI Readiness Pulse, Q2 2026 (v0.2A). Scores measure public evidence, not verified capability.



This is a ranking of records, not of companies.

A company reaches this table by publishing what it does with AI where it can be found.

One absent from it may be doing more, and saying less.



67.1

Average across these ten. Every ranked company here averages 44.2.



2

Of the ten rest on grade A evidence, the most traceable the index reads.



4 of 10

Of the ten are in professional services.

How Japan's numbers were built

Three steps produce every figure in this report. Most readers only ever see the third. The first two are where the honest limits live.

1. What we look at, and only what we look at

Every Japan score is built from public material: company sites, regulatory filings, job postings and third-party coverage. Nothing is inferred and nothing is self-reported, so no company can improve its position by telling us it deserves to. We search in English, so work published only in another language is undercounted, and we carry that as a known blind spot rather than correcting it quietly.

2. The threshold, where most companies stop

Of 346 companies screened in Japan, 29.8% leave too little public material to score at all, against 42.3% across the region. They are not ranked low. They are not ranked. They receive no score, no tier, and appear in no table in this report.

3. The score, and what it is a score of

The 243 companies that clear the threshold receive a visibility score, which measures how much of their AI work is publicly evident and never whether that work is any good. A company here can be genuinely capable and still score low, and the distance between the two is the finding. Tables here show 232 rows, because where one group is scored under several web domains only its strongest row is listed.



A missing score is not a low score.

It is the most common misreading of this index, which is why the threshold gets a step of its own above.

What this report does not claim

Printed with the findings rather than in an appendix at the back.

Known blind spots in this market

Japan carries 2 this cycle: evidence is gathered from English-language search, likely undercounting local-language activity; below the full three-source-block coverage standard. Each is a direction we know the number leans in. The country evidence behind the readiness score is graded B on source coverage of macro=1, enterprise=1.

A high score means published, not proven

A high visibility score means a company publishes a lot that can be found about its AI work. It does not mean the work is good, in production, or making money. A company here running serious AI quietly will score low here, and it should: this index reads the public record and nothing else.

What would change these numbers

Three things, all scheduled and all documented on page 21. Fuller source coverage per market, which is what currently sets the width of every readiness range. Validating the coverage weighting behind those ranges, which is a stopgap in this edition. And merging companies that operate several web domains, so one group cannot appear twice. Nothing here was audited inside a company, and no verified-readiness score is published this cycle. Next edition: Q3 2026.



103

Companies screened and left unranked. That is a missing record, not a low score.



68

Scored companies rest on grade C or D evidence and carry more uncertainty than the rest.




2

Known blind spots in this market, each published rather than quietly corrected.

What to do with a visibility score

A visibility score is something a company here can move, because every signal behind it is something it already controls. Three steps separate it from a record that shows what it is actually doing.



Start with the company's own row

Every scored company in Japan appears in the dataset with its five signal scores and an evidence grade. A grade C or D row is usually a company that has not said much, not a company that has not done much.



Then measure what this cannot see

This index reads the outside only. Leadership alignment, governance, data quality and workforce capability are invisible to it by design, and that is where the real answer sits.



Correct the record on purpose

Proof of capability is the region's weakest signal, because companies announce intent faster than they show work. Job posts, named leads and dated case studies move the score honestly.



None of this requires a new AI programme.
It requires publishing the one that already exists.



FIND YOUR OWN ROW

airapac.org/research

Dataset, method and corrections,
published in full.

About AIR APAC

An independent, vendor-neutral center for AI readiness in Asia Pacific. Ready people. Reliable processes. Trusted data.

The technology for AI adoption is mature. Most organisations are not. AIR APAC prepares organisations across people, process and data before they commit capital, and publishes the evidence base that informs that work.

Three principles

01

Neutrality

The Center is vendor-agnostic. It does not sell software and takes no commissions from technology vendors. Its role is to say whether an organisation is ready to adopt, not what it should buy.

02

Evidence

Opinion is insufficient for strategic decisions. This index is one output of a standing public-signal dataset on organisational maturity across APAC. Guidance is benchmarked against evidence, not theory.

03

Context

Frameworks written elsewhere do not always translate to Southeast Asia, North Asia or Oceania. The Center works inside the regulatory and institutional dynamics that define the region.

What the Center publishes



Research

This index, the per-market context benchmarks and the methodology behind both.



Programs

Certification and executive education on the human layer of AI adoption.



Advisory

Readiness assessment covering what public evidence structurally cannot see.

Research and advisory work rests on standing executive, academic, technology and advisory councils, none of which carry vendor sponsorship. Exeter Labs is the engineering arm behind the evidence infrastructure; readiness guidance never depends on buying its software.

Sources and notes

Provenance

Index	AIR APAC AI Readiness Pulse, Q2 2026.
Algorithm	ai_readiness_public_signal_v02b
Run	2026-Q2, generated 2026-08-01.
Method	docs/COUNTRY_READINESS_TRIANGULATION_METHOD_2026Q2.md and airapac.org/research/methodology .
Next edition	Q3 2026. Corrections received before that run are applied in it.
Coverage	5,820 companies screened region-wide, 3,359 of them scored, across 22 markets. 3,232 rows after merging multi-domain duplicates for display.
This market	346 companies screened in Japan, 243 scored and 103 not, 232 rows shown, 83 listed.

Method decisions and open limits

Declined. Derivative Oxford/IMF meta-indices. Conditionally open. Independent sources that fill Taiwan, Hong Kong or Macau; without them grade A is unreachable, but markets are not frozen at grade C. Measured, not corrected. None yet; recency and local-language effects remain deferred. What none of this fixes. Depth and evidence-density overlap still closes cross-market comparison.

Note

Scores in this report are a measure of publicly observable evidence. They are not an audit, not a certification, and not a statement about any organisation's internal capability, financial performance or prospects. An organisation may run substantial AI programmes and score low here if it does not publish them; the reverse is equally possible.

Ranks are region-wide and are given only to companies that clear the evidence threshold. Absence from this report is not evidence of absent activity. Evidence is gathered in English, so activity published only in another language is systematically undercounted, and that limit is published as a known blind spot rather than corrected silently.

The country readiness score (CERS in the dataset) is a separate measure from a company's visibility score (RPS) and is never added to it. Cross-market comparison is not licensed for this edition because sourcing depth and public evidence density are confounded.

AIR APAC does not accept technology vendor funding and takes no commission on any product referenced in this report. Corrections to any published row are accepted at airapac.org and applied in the following quarterly cycle.

How to cite

AIR APAC (2026). AI Readiness Pulse: Japan, Q2 2026. Center for AI Readiness in Asia Pacific. Retrieved from airapac.org/research.

Where the build differs from the written method

This page is for reviewers and is deliberately technical. Every build differs from its specification somewhere; these are the places this one does, recorded by the build itself and reproduced without editing.

01

coverage_quality (cov_q)

Specified. Section 15.1 lambda uses cov_q but does not define it.

Built. Proxy cov_q = $0.5 * \text{share}(\text{companies with evidence_confidence A/B}) + 0.5 * \min(1, n_{\text{effective}}/50)$. This stopgap silently drives lambda and therefore CERS; it is not validated. Replace when real coverage data exists.

02

services_class R/O/C/P

Specified. Section 7 requires multi-signal classification; warns name regex catches ~5/738.

Built. Implemented as NAME-ONLY word-boundary heuristic for precision. HIGH precision, LOW recall.

Consequence: the company-level services_class_factor penalty (O->0.5, R->0.0) rarely fires, so the company-level BPO correction relies primarily on (a) the evidence-based company_hiring_intent_factor (official-careers vs third-party job-board counts, which still fires for all rows) and (b) the country posting-bias factor (macro CRP-derived). The offshore-hub hiring dampening is therefore partly macro, not fully company-level. Disclosed per Section 0.1.

03

CERS precision

Specified. Section 12 says publish bands, not false precision.

Built. CERS published as integer plus a band spanning the k=75/150/300 sensitivity results. Precise sensitivity values are retained in the audit JSON.

04

evidence_confidence A/B/C-D

Specified. Section 12.1 needs official/third-party/job evidence grading.

Built. Heuristic from all per-dimension evidence URL host classification + dimension scores. Not validated.

NOT_BUILT - aer_score_or_null is NULL for all rows. The Salesforce/Access Partnership Business AI Readiness source now provides an enterprise/business block for 12 APAC economies, but company evidence confidence and services/entity classification remain heuristic and the full >=3-block CRP v0.2 coverage standard is not met. adjusted_rps (hiring-corrected) and cap_rules_applied are published as audit intermediates, not as company AER.

Japan, Q2 2026

Japan's record is current and it is cited. The 17 companies that also publish proof of capability average 62 of 100, on a method the rest can copy.



103

Companies here that buyers and investors cannot assess at all.



149

Scored but unlisted, and closest to moving up a tier.



17

Companies here already publishing proof of capability, averaging 62.

[READ THE FULL INDEX](#)

airapac.org/research

Nothing in this report rests on a source a reader cannot check.



The index

Every listed company, with its five signal scores and evidence grade.



The method

How each signal is weighted, where the threshold sits, and how the readiness score is built.



Corrections

Dispute a row any time. Applied in the Q3 2026 edition.



WORK WITH AIR APAC

airapac.org/scorecard

Readiness advisory and executive programs.
No software, and no vendor commissions.