

THE CRANE CONSULTANCY

BUSINESS INTELLIGENCE, PACED BY PROFIT

CRANE INDEX™ RESEARCH

UK Airports: The AI Visibility Report

Three machine reads of the 32 organisations in the UK Airports cohort, measured 2026-10-01. Can AI find you, can an agent act on you, does AI know who you are.

22	47	0	10
SITES READABLE	MEDIAN CRANE INDEX™	SITES AT 70+	UNMEASURED

Nobody in this cohort is ready. That is the opportunity.

Of the 32 sites approached, 10 could not be measured at all: 9 refused our reader outright, 0 gave it no response, and 1 serve pages that are empty until scripts run. Of the 22 sites our reader could read, not one reached 70 out of 100 on the composite Crane Index™. The three reads tell one story: access is broadly survivable (Agent median 63), meaning is patchy (AI Search median 54), and identity is the collapse (Representation median 12, with 22 of 22 sites below 55).

The strongest single separator between the top of this cohort and the bottom is not budget, brand or size. It is whether the site declares, in its own structured data, which organisation it is. 8 of the 10 highest-ranked sites here do; 0 of the 10 lowest-ranked do. That is a markup and governance exercise rather than a rebuild, which is why the league table overleaf should be read as unclaimed territory.


The three reads

AI Search Readiness. Median 54, range 21 to 80. Whether the engines that read live sites (AI Overviews, Perplexity, ChatGPT with search) can resolve the company and extract what it sells, to whom.

Agent Readiness. Median 63, range 34 to 87. Whether an automated agent is admitted at the edge and finds anything machine-operable when it arrives.

Brand Representation. Median 12, range 8 to 46. Whether a machine can confirm from the site itself which real-world organisation it is.

The picture in one chart



The story in one line: this cohort is led by Cornwall Airport Newquay on 65, while Liverpool John Lennon Airport sits on 23. Rank in this table is decided by structure, not size or spend, and the top-right quadrant, found *and* trusted, is empty. Whoever claims it first owns it.

What the identity gap actually is

Representation is not one judgement. It is four measured sub-scores, weighted and combined, and naming them turns a low number into a work list. The weights are the scanner's own; the medians are this cohort's.

SUB-SCORE	WEIGHT	MEDIAN	SCORING ZERO	WHAT EARNS THE POINTS
Identity	30%	0	14 of 22	An Organisation, Brand or Person declared in the site's own structured data, with a description and a depth field (knowsAbout, areaServed, founder, foundingDate); then that same name echoed in the page title or og:site_name, with a legal identifier (company number, VAT, LEI).
Disambiguation	30%	0	16 of 22	sameAs links to authoritative homes that resolve <i>and</i> name the organisation back (two score in full, one scores half, a link we cannot check counts half of one we can); and cover across kinds, a confirmed registry such as Companies House, Wikidata or OpenCorporates alongside a distinct profile.
Authority	25%	0	12 of 22	A named person behind the content, marked up with credential depth (job title, employer, expertise or profiles), plus specific figures and outbound sources a model can quote and check.
Rights	15%	50	0 of 22	Terms declared in the markup: a schema license or copyrightHolder, a TDM reservation, or a noai directive. This check can be met in part or in full but cannot be failed, so 50 is its floor, not a result.

Rights cannot be failed, only met in part or in full, so a median of 50 on that row is the floor of the scale rather than a result. Read the identity and disambiguation rows first.

What a zero means for a business

Zero on identity (14 of 22 sites here). There is no Organisation, Brand or Person declared in the site's structured data at all, so there is no name for a machine to check against the page it is reading. The site tells a human visitor who it is and tells a machine nothing. Everything downstream rests on this: a model cannot confirm an identity that was never declared.

Zero on disambiguation (16 of 22 sites here). Nothing on the site links the organisation to an authoritative home that both resolves and names it back. Either no such links are declared, or the ones declared did not resolve or did not mention the organisation. A machine can read a name and still cannot tell which organisation of that name it has found, which matters most in a sector where several share one.

The same four reads, two sites

	IDENTITY	DISAMBIGUATION	AUTHORITY	RIGHTS	REPRESENTATION
Cornwall Airport Newquay cornwallairportnewquay.com · ranked 1	50	0	60	50	38
Southampton Airport southamptonairport.com · ranked 17	0	0	20	50	13

Why the pair. Cornwall Airport Newquay leads this cohort. Southampton Airport is the ordinary case, the site whose representation score sits closest to the cohort median of 12. It is ranked 17 of 22 overall. The gap between the two rows is not effort or budget: it is markup that one site publishes and the other does not. These are scores for what each site declares about itself to a machine, and say nothing about either organisation.

Every site our reader could read

Rank is by the composite Crane Index™. A low score says the site is hard for a machine to read, not that the business is anything other than excellent at what it does. Sites that refused or could not serve an automated reader are excluded as unmeasured, never scored zero.

#	COMPANY	CRANE INDEX™	AI SEARCH	AGENT	REPRESENTATION
1	Cornwall Airport Newquay cornwallairportnewquay.com	65	80	76	38
2	East Midlands Airport eastmidlandsairport.com	56	60	63	46
3	Leeds Bradford Airport leedsbradfordairport.co.uk	56	64	62	41
4	London Southend Airport londonouthendairport.com	56	72	87	8
5	Norwich Airport norwichairport.co.uk	56	65	79	25
6	London Stansted Airport stanstedairport.com	55	60	63	43
7	Manchester Airport manchesterairport.co.uk	54	55	63	43
8	Heathrow Airport heathrow.com	53	58	69	32
9	London Gatwick Airport gatwickairport.com	52	56	62	37
10	Glasgow Prestwick Airport glasgowprestwick.com	51	57	85	10
11	Exeter Airport exeter-airport.co.uk	48	57	79	8
12	Teesside International Airport teessideinternational.com	46	53	77	8
13	London Luton Airport london-luton.co.uk	41	48	68	8
14	Edinburgh Airport edinburghairport.com	39	46	62	8
15	Aberdeen International Airport aberdeenairport.com	38	44	59	10
16	Inverness Airport hial.co.uk/inverness-airport	37	45	56	10
17	Southampton Airport southamptonairport.com	37	44	53	13
18	Kirkwall Airport hial.co.uk/kirkwall-airport	36	45	56	8
19	Stornoway Airport hial.co.uk/homepage/9/stornoway-airport	36	44	56	8

#	COMPANY	CRANE INDEX™	AI SEARCH	AGENT	REPRESENTATION
20	Sumburgh Airport hial.co.uk/sumburgh-airport	36	45	56	8
21	Humberside Airport humbersideairport.com	23	21	34	15
22	Liverpool John Lennon Airport liverpoolairport.com	23	21	34	15

The 10 sites our reader could not measure

Unmeasured, not failed. Our reader is one unfamiliar automated visitor, and how a site treated it is all that was measured. No score is asserted for any site on this page.

What was measured, posture by posture. Of the 10 sites here, 9 refused our reader, 0 gave no response and 1 served a page that is empty until scripts run.

Refused the reader (9). The edge answered our request with a refusal. That is a measurement of how one unfamiliar visitor was treated and nothing further: a site that admits crawlers by name may well serve GPTBot or Googlebot while refusing us, so a refusal here is not evidence that the site is invisible to any named engine, and it is not reported as one. What it does establish is that a rule at the edge is deciding which machines are served, which is worth knowing deliberately.

Script-only page (1). The page returns almost nothing until a browser executes its scripts. This one does generalise: a page that is empty without JavaScript is empty for every fetcher that does not run it, ours included, and many AI fetchers and shopping agents do not run it. Here the commercial reading is direct, because the page has nothing to serve any of them.

Birmingham Airport · refused the reader	Newcastle International Airport · refused the reader	Bournemouth Airport · refused the reader
Bristol Airport · refused the reader		City of Derry Airport · refused the reader
Glasgow Airport · refused the reader	London City Airport · refused the reader	Cardiff Airport · script-only page
Belfast International Airport · refused the reader	George Best Belfast City Airport · refused the reader	

Method

Each organisation's primary public site was loaded and scored by the three deterministic Crane Index™ scanners on 2026-10-01. The constituent list is pinned beside the archived results, exclusions are rule-based and inspectable, and the full per-check result set is retained so the cohort can be re-scanned and the movement published. Every score can be re-measured in about a minute at thecraneconsultancy.com/ai-readiness/.

How this cohort was drawn

The list of 32 organisations is assembled from public sources, sanity-checked by a person, and pinned beside the archived results so the composition can be audited rather than taken on trust. Anyone can ask for a domain to be added; it joins only after a human review, against one bar, a real trading organisation of substance in this industry, and requested-in constituents are recorded as such. Two site types are excluded uniformly, by type and before any score is seen, because the measure is not a fair test of them: search-engine homepages and chat-box homepages. Nothing is ever removed for scoring badly. So this is a defined and inspectable list of 32 organisations. It is not a ranking of national importance, and it is not a census of the industry.

Honest limits

22 scored sites is a sample, not a census, and a single scan is a snapshot of engines and estates that shift month to month. Our reader is one unfamiliar bot, so a site that refused it is reported as unmeasured, never scored zero and never described as unreadable to AI, because what a named engine is served was not measured here. The medians on these pages therefore describe the 22 sites our reader could read; how the 10 it could not would have scored is unknown, and this report does not assume they would have scored worse.

The three moves to make first

1. Ship valid Organization and Product structured data (JSON-LD) on every commercial template, matched to the visible copy, one legal name rendered identically everywhere, with sameAs disambiguation. 2. Make your bot posture a decision rather than an inherited default: know which automated readers your edge admits and which it turns away, name the engines that matter to your market, and record the choice where the next person will find it. 3. Put a legibility gate in the deploy pipeline and give one senior owner the score, because estates drift the moment legibility has no owner.