

THE CRANE CONSULTANCY

BUSINESS INTELLIGENCE, PACED BY PROFIT

CRANE INDEX™ RESEARCH

The FTSE 100: The AI Visibility Report

Three machine reads of the 79 organisations in the The FTSE 100 cohort, measured 2026-09-01. Can AI find you, can an agent act on you, does AI know who you are.

40	43	0	39
SITES READABLE	MEDIAN CRANE INDEX™	SITES AT 70+	UNMEASURED

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

Of the 79 sites approached, 39 could not be measured at all: 23 refused our reader outright, 10 gave it no response, and 6 serve pages that are empty until scripts run. Of the 40 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 64), meaning is patchy (AI Search median 48), and identity is the collapse (Representation median 14, with 40 of 40 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. 10 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.

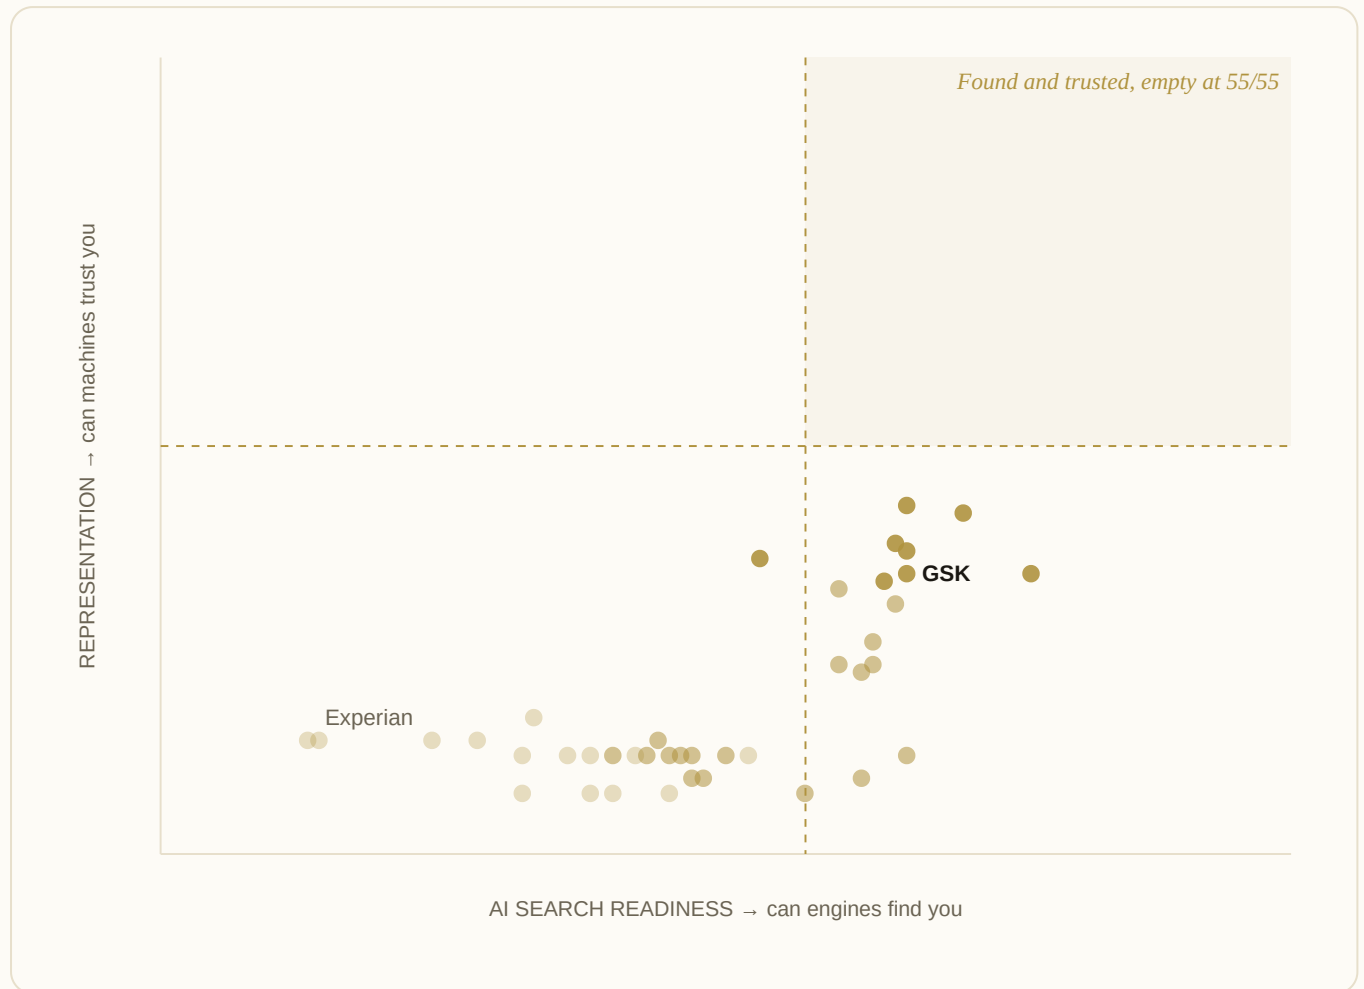
The three reads

AI Search Readiness. Median 48, range 13 to 77. 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 64, range 30 to 89. Whether an automated agent is admitted at the edge and finds anything machine-operable when it arrives.

Brand Representation. Median 14, 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 GSK on 64, while Experian sits on 21. 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	26 of 40	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	34 of 40	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%	20	12 of 40	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 40	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 (26 of 40 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 (34 of 40 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
GSK gsk.com/en-gb · ranked 1	50	30	20	50	37
Marks & Spencer marksandspencer.com · ranked 14	0	0	20	50	13

Why the pair. GSK leads this cohort. Marks & Spencer is the ordinary case, the site whose representation score sits closest to the cohort median of 14. It is ranked 14 of 40 overall, in the stronger half of the table. 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	GSK gsk.com/en-gb	64	66	89	37
2	Taylor Wimpey taylorwimpey.co.uk	63	65	83	41
3	NatWest Group natwestgroup.com	62	71	69	45
4	Whitbread whitbread.co.uk	62	66	79	40
5	Reckitt reckitt.com	59	77	63	37
6	Persimmon persimmonhomes.com	58	66	62	46
7	Phoenix Group standardlifeplc.com	57	64	71	36
8	Weir Group global.weir	56	53	76	39
9	Schroders schroders.com/en/global/individual	54	65	65	33
10	British American Tobacco bat.com	53	63	68	28
11	Kingfisher kingfisher.com	52	60	62	35
12	Ashtead Group ir.sunbeltrentals.com	49	63	58	25
13	Croda International croda.com/en-gb	49	60	62	25
14	Marks & Spencer marksandspencer.com	48	45	87	13
15	HSBC Holdings hsbc.com	47	62	68	10
16	Legal & General legalandgeneral.com	47	66	62	13
17	Convatec convatecgroup.com	46	62	53	24
18	JD Sports Fashion jdplc.com	44	50	68	13
19	London Stock Exchange Group lseg.com/en	44	57	68	8

#	COMPANY	CRANE INDEX™	AI SEARCH	AGENT	REPRESENTATION
20	Mondi mondigroup.com	44	47	71	13
21	Smiths Group smiths.com	42	47	68	10
22	United Utilities unitedutilities.com	42	48	68	10
23	Rightmove rightmove.co.uk	41	44	65	15
24	Smith & Nephew smith-nephew.com/en	41	43	68	13
25	Spirax Group spiraxgroup.com	41	40	69	13
26	Rio Tinto riotinto.com	40	46	62	13
27	International Consolidated Airlines Group iairgroup.com	39	40	68	8
28	Intertek Group intertek.com	39	45	63	8
29	Prudential prudentialplc.com/en	39	42	62	13
30	SSE sse.com	39	36	68	13
31	Pearson pearson.com/en-us	38	38	62	13
32	Glencore glencore.com	36	32	62	13
33	Landsec landsec.com/en	36	33	56	18
34	WPP wpp.com/en	35	38	58	8
35	Beazley beazley.com	33	28	55	15
36	Severn Trent severntrent.com	33	32	59	8
37	Barclays home.barclays	32	52	30	13
38	Lloyds Banking Group lloydsbankinggroup.com	24	24	34	15
39	BAE Systems baesystems.com	21	14	34	15
40	Experian experianplc.com	21	13	34	15

The 39 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 39 sites here, 23 refused our reader, 10 gave no response and 6 served a page that is empty until scripts run.

Refused the reader (23). 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.

No response (10). Nothing came back inside our timeout. That could be an outage on the day, a network rule or a slow origin, and the scan cannot tell which, so nothing is concluded about the site from it.

Script-only page (6). 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.

AstraZeneca · refused the reader	InterContinental Hotels Group · refused the reader	Next · no response
Unilever · refused the reader		Segro · no response
BP · refused the reader	Informa · refused the reader	Antofagasta · no response
Diageo · refused the reader	Rentokil Initial · refused the reader	Fresnillo · no response
National Grid · refused the reader	Sage Group · refused the reader	Berkeley Group · no response
Rolls-Royce Holdings · refused the reader	St James's Place · refused the reader	Admiral Group · no response
Aviva · refused the reader	J Sainsbury · refused the reader	Standard Chartered · script-only page
Vodafone Group · refused the reader	DCC · refused the reader	Compass Group · script-only page
Tesco · refused the reader	Melrose Industries · refused the reader	Associated British Foods · script-only page
3i Group · refused the reader	Vistry Group · refused the reader	Coca-Cola HBC · script-only page
Anglo American · refused the reader	Shell · no response	Haleon · script-only page
Auto Trader Group · refused the reader	RELX · no response	Imperial Brands · script-only page
Centrica · refused the reader	Bunzl · no response	
Halma · refused the reader	M&G · no response	

Method

Each organisation's primary public site was loaded and scored by the three deterministic Crane Index™ scanners on 2026-09-01. 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 79 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 79 organisations. It is not a ranking of national importance, and it is not a census of the industry.

Honest limits

40 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 40 sites our reader could read; how the 39 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.