2026 AGENT READINESS INDEX

What the public web reveals to agents

A bounded, unauthenticated structural assessment of 120 companies across ten industries.

120	71	49
Companies attempted	**Usable projections**	**Not assessed**
All frozen targets	59.17% of 120	40.83% of 120

The measured evidence gap

25
Median confirmed points / 100



50
Median weighted coverage / 100



Both medians are calculated only over the 71 usable provisional results. A confirmed point is supported by observed evidence; a point outside assessed coverage remains unresolved or unmeasured. These numbers are not estimates of an organization’s complete AI capabilities.

Three findings

1. Access matters. Forty-nine companies could not be assessed, while 17 others retained usable evidence despite a blocked measurement classification.

2. Discoverability is not a machine interface. Only 13 of 71 assessed companies had a verified MCP or OpenAPI contract in this scan.

3. Declared authorization is not enforcement. Eight of 71 received declared-contract credit. No operational or enforcement tests were performed.

Scope: purposive, balanced cohort; not a random or representative sample of the web. Scan date: 28 November 2026. All 120 workers completed without recorded errors or timeouts.

ACCESS & AVAILABILITY

Where evidence was obtainable

Measurement health is separate from whether the scanner can support a provisional assessment.

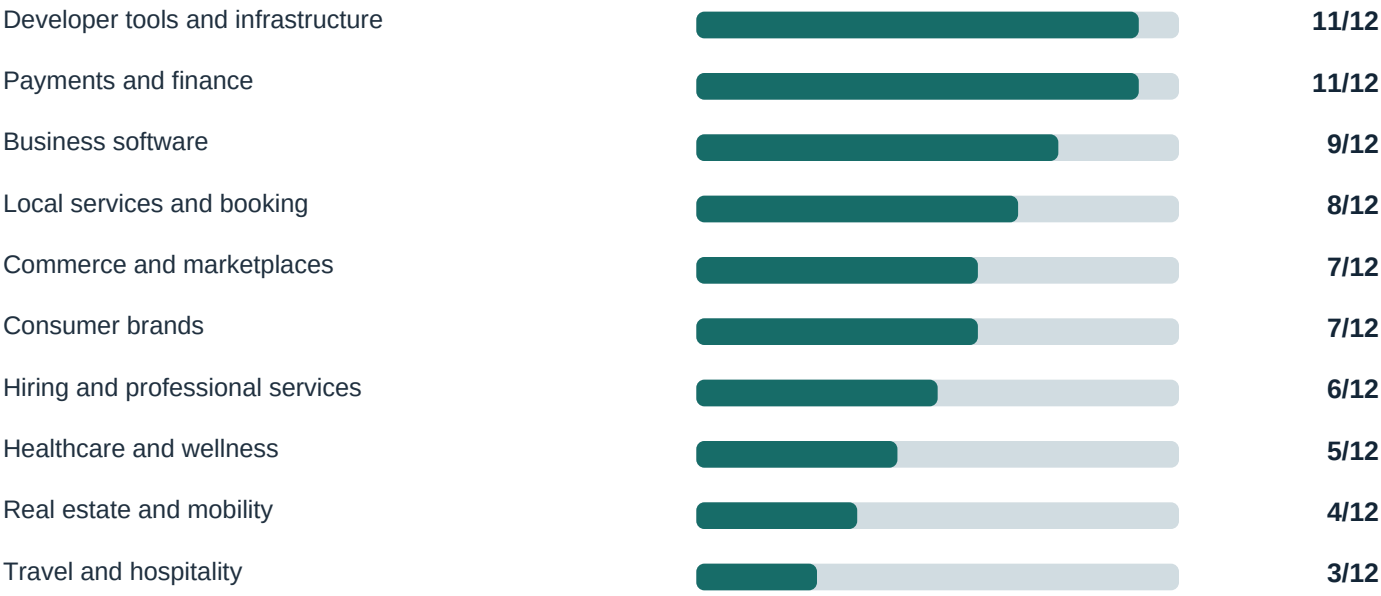

Each sector contains twelve preselected companies. These rates reflect this run’s access and evidence availability, not relative industry readiness or demand.

Measurement	Provisional	Not assessed	Total
Complete	27	0	27
Partial	27	1	28
Blocked	17	32	49
Challenged	0	2	2
Rate limited	0	3	3
Network failed	0	11	11

Homepage health: 71 complete, 31 blocked, 2 challenged, 2 partial, 3 rate-limited, and 11 network failed (120 total). HTTP 403 was observed on 31 homepages. Not all blocked measurements arose from blocked homepages: secondary probes also encounter scope and media restrictions.

SCORING & UNCERTAINTY

Confirmed points are not coverage

The scoring model allocates 100 weighted points across four dimensions.

Dimension	Weight	Evidence basis
Machine interfaces	40	Verified MCP or OpenAPI contract
Discoverability	25	Catalog, guidance and site index
Semantic understanding	20	Entity and form structure
Access governance	15	Crawler policy and declared contract

Confirmed points sum evidence-supported credit. Weighted assessment coverage sums assessed weights, including resolved findings without credit (such as scoped absence or not-applicable results). Documented but unverified, unresolved and unmeasured evidence do not automatically earn points. Coverage is not a confidence interval, probability or percent of all site pages inspected.

Confirmed points

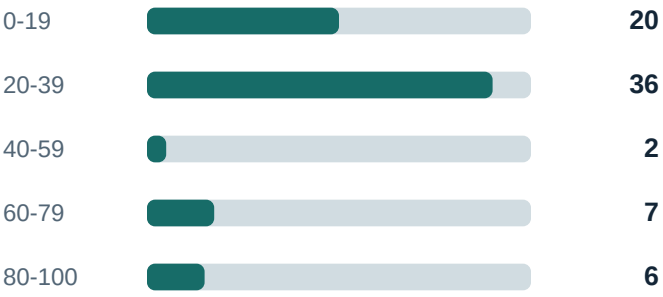

n = 71 provisional assessments

Weighted coverage

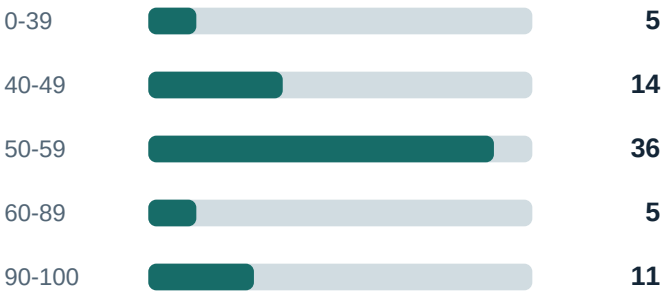

n = 71 provisional assessments

56/71 (78.9%) had fewer than 40 confirmed points. **55/71 (77.5%)** had 50 or fewer assessed points. Four had full weighted coverage, but three of those still had PARTIAL measurement health. Coverage and scan health answer different questions.

Do not infer a missing capability. Three provisional results had zero confirmed points; 49 results were not assessed and retain null values. Neither result establishes a company-wide lack of agent support.

PROTOCOLS & AUTHORIZATION

Contracts are not operational proof

Company-level protocol states among the 71 usable provisional assessments.

Protocol	Verified / shape	Documented	Other states
MCP	6 verified	13 unverified	37 absent in scope; 15 unresolved
OpenAPI	10 verified	-	25 absent; 31 unresolved; 5 invalid
ARD	7 shapes	-	55 absent; 5 unresolved; 4 invalid
A2A	2 shapes	-	59 absent; 6 unresolved; 4 invalid
CLI	-	10 unverified	53 absent in scope; 8 unresolved
WebMCP	-	-	71 unmeasured

MCP and OpenAPI counts overlap: 13 distinct companies had at least one verified contract. “Absent” and “invalid” here refer to the bounded probe scope or returned evidence, not a claim about the entire organization. ARD and A2A shape observations are not verified live services.

Declared authorization evidence

Observation	Companies	Denominator
OpenAPI declared contract observed	7	71
MCP resource binding observed	5	71
MCP metadata reference only	1	71
Awarded declared-contract credit	8	71

The credit rule awards up to 10 points for qualifying declared authorization evidence, not an enforcement test. The OpenAPI and MCP evidence counts overlap. Some declared OpenAPI observations do not qualify for credit; e.g., ClickHouse’s authorization-contract capability remained unresolved.

NOT TESTED: authenticated access, tool invocation, business-task completion, owner authorization, and runtime enforcement. Active MCP initialization was disabled throughout this benchmark. A robots.txt document is evidence of crawler policy, not proof of authentication or permission to execute actions.

EVIDENCE IN CONTEXT

Three instructive cases

Selected to illustrate evidence and scope, not to rank vendors. Scores are confirmed points / weighted coverage.

Resend 90 / 100

Measurement: COMPLETE | Projection: provisional

Verified MCP and OpenAPI contracts, a catalog shape, guidance, sitemap, entity structure and declared OpenAPI authorization. Form structure was not applicable (10 assessed points without credit). This is strong evidence within the rubric, not a test of production behavior.

Recorded sources: MCP manifest | OpenAPI document

ClickHouse 80 / 90

Measurement: COMPLETE | Projection: provisional

The scan verified MCP and OpenAPI contracts, including the explicitly linked external OpenAPI document. A declared OpenAPI authorization observation was recorded, but the authorization-contract capability remained unresolved and earned no credit.

Recorded sources: MCP document | OpenAPI document

GitHub 40 / 50

Measurement: COMPLETE | Projection: provisional

The scanner found catalog structure, guidance, a sitemap and documented MCP references, but did not verify an MCP or OpenAPI contract. The result describes bounded discovery, not the absence of GitHub interfaces.

Recorded sources: Catalog evidence | MCP documentation

Contrast: REI returned HTTP 403, so its score and coverage are null. **Attio** had a blocked secondary measurement but still retained 90 confirmed points and 90 assessed points. Neither case belongs in a simple best-to-worst ranking.

METHODOLOGY & REPRODUCIBILITY

Reproducible measurement, bounded claims

Source: Arachne's 28 November 2026 integrated benchmark dataset and companion aggregate.

Design. Ten editorial sectors, 12 companies each. Entry URLs selected before scanning; domains were not independently preflighted. All 120 were attempted once, with one worker, a five-second inter-site delay, a 180-second per-worker timeout and no retry policy. Run window: 09:25:53–10:00:22 UTC. Public, passive structural discovery only; no login, purchase, task execution or active MCP initialization.

Scope. The run directly exercised static scanner projection, not the hosted frontend, API quotas or lead funnel. Bounded acquisition, media validation, namespace policies and deadlines constrain what is observed. DNS namespace relationships alone do not establish publisher authority. Interfaces may exist beyond the entry URLs, documented links or allowed acquisition scope. Results reflect a single measurement window and are subject to access-related selection bias.

Exact provenance

Commit
`d4c8bb81fbcdceea90f25b5529a33bf1d2b3dfc1`

Tree
`a313b669a0b31731b18768f61a4481cea6983d8e`

Scanner blob
`712d482c1ae6ec09cf4b62a6f35f4316b6beccb2`

Roster SHA-256
`bf498eb87eae26482a47fe1606a18daba6d37c96adb05c4f8476b0e99fba9b42`

Runner SHA-256
`724e837836d02b33a6dc63b12b21fc9828302d26a8c00d045d9b6c2999d42878`

Diagnostic SHA-256
`d5786b2f98bca4249679526c5d903fdb48ac1cb8c36fe5a67d8fb3c172e03dbd`

Dependency inventory SHA-256
`015ded833ba99e5be16893ba70a0e649d55bb17b17c0c63e4d962ade5d1063c6`

Dataset SHA-256
`ffd920ba46427a8edcced0bcbfd028d7e2d5ed8bed5883f1690a5b274b88b5c3`

Aggregate SHA-256
`d0aea070b525baf9d10d05a46a69cbd822ab75cf9ddd3938ed431d727a0eba8c`

Scanner: arachne-readiness-v3 · Scoring: arachne-readiness-v4 · Projection: arachne-unified-governance-shadow-v1 · Protocol snapshot: 2026-08-29 · Python 3.11.9 / Windows.

Validation: supplied file hashes matched; 120 unique rows matched the supplied roster; major counts, distributions, dimension totals and protocol states were independently recomputed. This is an audit of the supplied artifacts, not a fresh crawl or independent certification of each endpoint.