

Registry discovery refresh

October 12, 2026 | New searches, scoped additions and outstanding candidate reviews

October 12, 2026. This is a new discovery wave under the September Core relevance rules. Historical search and relevance audits retain their original dates and denominators.

Results

62 new Core capability records were added after static primary-documentation review, bringing the website catalogue from 4,218 to **4,280** records. Records span MCP servers, agent skills, components and workflows; these are not counts of unique MCP servers.

The reconciled primary-review queue contains 172 candidate identities: 62 included, 17 excluded, 88 held for further review, and 5 existing or reconciled identities. Held discovery candidates are outside the catalogue. They are separate from the one historical Borderline record and one historical implementation hold. The 2,008 archived catalogue records remain unchanged.

New IDs run from RWE-6914 to RWE-6975. New records retain retrieval timestamps, README blob identifiers, source-text hashes, scoped capabilities and limitations in the maintained JSON. Their source and relevance reviews are dated October 12. Historical cards were not relabelled as newly reviewed.

Sources actually searched

Source	Completed retrieval	Boundary
Official MCP Registry	216 pages; 21,514 update observations; terminal cursor reached	Latest-version updates since September 7, including deletion metadata. This is a delta feed, not a fresh full baseline harvest.
GitHub repositories	64 queries; 64 result pages; 432 overlapping observations	32 health/method concepts paired with MCP and skill signals in names, descriptions and topics; public non-fork repositories pushed since September 7. All queries reached short/empty terminal pages without reported incomplete results.
Smithery	32 queries; 64 pages; 4,616 overlapping observations	Current ranked vocabulary search, paged to its reported/observed end. Broad matches were screened locally. Live totals changed between requests.
Indexed directories	20 queries across named directories	Ranked indexed results from Glama, PulseMCP, MCP.so, MCP Market, MCP Repository and LobeHub; no full-directory inventory claim.
Curated healthcare list	Current README retrieved with blob identity	Awesome Healthcare MCP Servers used as a discovery source; its compliance and clinical ratings are not adopted.

Queries, page URLs, returned counts and primary-review dispositions are retained in the [structured refresh audit](#). Search observations overlap across terms, pages and sources; adding them does not produce a count of unique projects.

Review approach

Automatic metadata screens separated potential health/method tools from unrelated matches. This was not individual source review. Promoted GitHub candidates were reconciled with existing canonical repository URLs and their primary documentation was retained. Only explicitly adjudicated records were published. Unknown identities, insufficient documentation and unresolved component/implementation questions remain held, with reasons in the audit.

Core requires a documented health-research, clinical-data, public-health, evidence-synthesis or specific statistical-method function. Mixed collections are limited to the qualifying function named on each card. Generic productivity, molecular-only uses and possible downstream health applications do not independently qualify. Mock-only FHIR operations, registry self-listings and biological MCP-1 homonyms were excluded in this wave.

The new review examined documentation and repository metadata. It did not inspect all implementation files, install candidate projects or call their tools. Author test counts, compliance statements, model performance and benchmark results are not independent registry validation. Synthetic examples are explicitly labelled; they are not real clinical data.

Examples of added scope

- Survey-weighted NHANES analysis: [NHANES MCP](#). This is a Black Swan Causal Labs product and retains the ownership marker.
- Exploratory adverse-event methods: [FAERS MCP](#), with reporting-bias and non-causality limitations.
- International trial registries: [ANZCTR lookup](#), [CTIS](#), [ISRCTN](#) and [ICTRP](#), each limited to its named adapter.
- Evidence screening and verification: [TypeSafe screening](#), [Pho Research Tools](#) and [Librarian Search](#).
- Standards and measure metadata: [CDISC Library MCP](#) and [Open Quality](#).

Remaining coverage limits

This wave reruns bounded discovery; it does not claim exhaustive coverage, the original Cartesian matrix, or two successive no-novelty rounds. 88 candidate identities still need first-party identity, distinctness or component-level follow-up. No new GitHub code-search lane, full README-index matrix, full Docker catalogue audit or whole-catalogue capability re-review was executed. Unchanged older projects without matching metadata or directory results can still be missed. Some hosted endpoints may be aliases of known tools that require further identity reconciliation.

The official update feed is live rather than a transactionally frozen snapshot. Deletions and metadata changes do not by themselves establish the availability of an underlying repository or hosted service. The overlap date reduces boundary omissions but does not prove that every old record has been rechecked.