MCP SERVER | NO CODE | CLOUD HOSTED

Stanford CrossRef MCP

Track research impact and funding from any paper.

Stanford CrossRef MCP connects your AI agent to global academic metadata. It lets you resolve Digital Object Identifiers (DOIs) and retrieve full publication details, citation counts, and funding information for millions of scholarly works. Use this MCP to quickly assess research impact, track institutional output, or build comprehensive bibliographies without leaving your chat window.

A+ Quality Score 98.33/100

crossref | doi | metadata | journals | publishers | funders

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/stanford-crossref

SCAN TO VIEW

The connectivity layer between AI and the world's software.

Vinkius sits between AI and every application. All communication passes through Vinkius Cloud via the Model Context Protocol (MCP) — with governance, observability, and security at every layer.

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the connectivity layer where AI connects to the software your business already runs. We handle the hosting, the security, the credentials, the uptime — you get agents that actually do things.

We operate the world's largest managed MCP catalog. Major SaaS platforms, CRMs, databases, and cloud providers — running, monitored, production-ready. This MCP server is hosted and maintained by the Vinkius Cloud for AI Agents.

> *The agent doesn't manage credentials, doesn't manage uptime, doesn't manage security. Vinkius does.*
>
> — Architecture principle

Four Pillars of the Vinkius Runtime

01 — Security by design

Credentials stay encrypted at rest via AES-256. The AI agent never touches raw keys — they're injected into a sandboxed V8 isolate at runtime. Actions are logged, and connections have an emergency kill switch.

02 — Built on MCP Fusion

This MCP server was built with **MCP Fusion**, the open-source framework (Apache 2.0) that powers the entire Vinkius catalog. Schema-as-firewall strips undeclared fields, compiled PII redaction runs at zero overhead, and cryptographic lockfiles produce git-diffable audit trails.

03 — Deterministic observability

Eight immutable metrics per endpoint: request volume, p95 latency, error rate, active connections, cost attribution. A live payload feed logs every tool call with mutation detection.

04 — Autonomous operations

Servers are deployed, monitored, and patched autonomously. New capabilities and security patches ship weekly. Zero-downtime deployments ensure continuous availability across all managed MCP servers.

AES-256	Ed25519	24h TTL	V8 Isolate
Encryption at rest	PKI vault signatures	Ephemeral session keys	Sandboxed execution

One Token. Instant Access.

Every MCP server on Vinkius is accessed through a **Connection Token**. Tokens are generated in the cloud dashboard and produce a unique MCP endpoint URL. Paste this URL into any MCP-compatible client — no SDK required.

A single token can serve **multiple AI clients simultaneously**, or you can issue separate tokens per client for granular access control. Each token tracks its own request count, last activity timestamp, and can be individually enabled or revoked.

MCP ENDPOINT

```
https://edge.vinkius.com/{token}/mcp
```

Claude · Cursor · VS Code · Windsurf · Grok · Gemini

Security Is the Architecture

Security in Vinkius is not a feature — it's the foundation of the runtime. The gateway enforces multiple independent protection layers between AI agents and third-party APIs.

01 — Ed25519 PKI Vault

Every workspace has an Ed25519 Master Key. Session keys are generated ephemerally (24h TTL) and signed by the Master Key. Credentials never leave the vault boundary.

02 — V8 Isolate Sandboxing

Tool code runs inside isolated-vm V8 isolates with 64 MB memory caps and per-request timeouts. No filesystem access, no network access except through the SSRF-guarded fetch bridge.

03 — SSRF Guard

All outbound HTTP requests are DNS-resolved and validated before execution. Private IP ranges (10.x, 172.16-31.x, 192.168.x, AWS metadata 169.254.x) are blocked at the network layer.

04 — DLP & PII Redaction

A ResponseGuard pipeline intercepts every tool response. Configurable redaction patterns strip sensitive fields (emails, SSNs, card numbers) before data reaches the AI agent.

05 — Cryptographic Audit Trail

Every request is signed into a SHA-256 hash chain with Ed25519 signatures. Events form a tamper-proof, SIEM-exportable forensic record.

06 — Honeytoken Trap System

Phantom credentials are injected into isolated environments. If a honeytoken is used outside Vinkius infrastructure, the server is quarantined instantly.

Emergency Kill Switch

EU AI Act Art. 14(1) Compliant

The kill switch is an **emergency halt** mechanism — not a simple toggle. When triggered, it executes three actions atomically:

01 — Server deactivated

The MCP server is immediately taken offline across the entire cluster.

02 — All tokens revoked

Every connection token is invalidated. Total lockout — reconnection blocked until new tokens are issued.

03 — WebSocket connections killed

Active connections terminated via Redis pubsub broadcast. Propagates to every runtime node in the cluster.

Full Visibility. Zero Guesswork.

The Vinkius cloud dashboard includes a full MCP Governance suite — real-time analytics and security controls for production AI operations.

Control Plane

KPI dashboard with request volume, latency, success rate, token consumption, and AI-generated operational briefings.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Incident Log

Error trends, failure rates, status-code breakdowns, and forensic audit trail access.

> Get started at **cloud.vinkius.com** — connect your AI agent in under 60 seconds.

Stanford CrossRef MCP

16 tools available
Cloud-hosted on Vinkius

This MCP gives your AI client access to the core infrastructure of academic publishing. Instead of visiting journal websites or using multiple databases, you connect once through Vinkius and gain immediate access to millions of scholarly records. You can resolve any DOI into complete bibliographic metadata, instantly identifying authors, journals, and publication dates. Need to know how influential a paper is? You can retrieve citation counts directly. Want to track where the money comes from? The MCP finds works funded by specific organizations like the NIH or NSF. It also lets you search for publications based on institutional affiliations or even find preprints before they hit a journal. This capability fundamentally changes research analysis, allowing your agent to build complex bibliographies and evaluate entire fields of study using structured data.

Core Capabilities

01 — Resolving Full Paper Metadata

Provide any DOI and receive immediate details including the title, authors, journal name, publication date, and subject areas.

02 — Assessing Research Influence

Determine how often a paper has been cited by other researchers or how many sources it cites itself.

03 — Mapping Funding Sources

Search and retrieve lists of academic works tied to specific funding organizations, like the National Science Foundation.

04 — Discovering Publication Details

Look up journals or publishers by ISSN or name to check their quality metrics, total output count, and coverage.

05 — Finding Works by People or Institutions

Search for publications using a researcher's universal ORCID identifier or an institution's name like Harvard University.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/stanford-crossref** — connect your AI agent in three steps.

01 Give your agent a query, such as 'Find the citation count for DOI X' or 'List all works by NIH'.

02 The MCP sends that request to CrossRef and pulls structured data (titles, dates, counts) from their massive index.

03 Your AI client receives clean, actionable metadata that it can then summarize, format, and present back to you.

The bottom line is: your agent treats the entire academic publishing record like a searchable database accessible via conversation.

Built For

This MCP is essential for academics who deal with massive amounts of publication data. It's for researchers tired of switching between citation databases and institutional repositories, and for librarians needing to evaluate journal quality at scale.

Academic Researcher

Using the MCP, they can resolve a paper's full record with `resolve_doi` and instantly check its influence using `get_citations_count`.

Librarian / Subject Specialist

They use this to evaluate journal quality by running `search_journals` or tracking which publishers are gaining traction via `search_publishers`.

Research Administrator

These users track institutional output and funding trends, using tools like `get_funder_works` to prove grant impact.

What Changes When You Connect

01 You instantly assess a paper's influence by running `get_citations_count` on its DOI, seeing immediately how many other researchers have used it. This is much faster than manual database checks.

02 Map the financial backbone of research using `get_funder_works`. Instead of guessing funding sources, you can pull structured data showing which grants funded a specific body of work.

03 Build robust bibliographies effortlessly. Use `get_reference_list` to extract every source cited in a paper; it's perfect for literature reviews and deep dives into academic foundations.

04 Evaluate the quality of journals or publishers using `search_journals` or `get_publisher`. You get metrics like metadata coverage percentages, helping you select reliable sources quickly.

05 Track researcher output efficiently. If you know an ORCID ID, use `search_by_orcid` to pull every single publication associated with that person's career record.

Real-World Applications

Writing a Literature Review on Gene Therapy

A researcher needs to know all high-impact papers funded by the NIH. Instead of running multiple database queries, they ask their agent to use `get_funder_works` with the correct funder ID and then run `resolve_doi` on the resulting list to get full titles, saving hours.

Evaluating a Journal for Publication

A librarian needs to recommend a new journal. They ask the agent to use `search_journals`, which returns not just the title and ISSN, but also total DOI counts and quality metrics needed to make an informed recommendation.

Attributing Work to a Specific Lab

A research administrator needs to confirm all publications by a specific professor. They use `search_by_orcid` with the person's unique identifier, getting a complete and verifiable publication list across multiple journals.

Patterns to Avoid

Treating DOIs as simple URLs

✗ AVOID

Copying a DOI (e.g., 10.xxxx/abc) into a browser and getting only an abstract page, forcing you to manually find the citation count or full reference list.

✓ INSTEAD

Use your agent to execute `resolve_doi` on the DOI. This forces the MCP to return structured data, giving you immediate access to authors, dates, and metrics like `get_citations_count`.

Searching journals in silos

✗ AVOID

Checking JSTOR for one journal and PubMed for another, leading to fragmented or incomplete views of a topic's output.

✓ INSTEAD

Use `search_works` with advanced filters (like filtering by 'type:journal-article') across the entire 150M+ DOI index. This gives you a unified view of multiple platforms.

Ignoring preprints

✗ AVOID

Only searching established, peer-reviewed journals, meaning critical, cutting-edge research from bioRxiv or SSRN is missed.

✓ INSTEAD

Use `search_preprints` to include the latest, non-peer-reviewed findings. This ensures your literature review is current and comprehensive.

The Right Fit

Use this MCP if you need verifiable, structured academic data that spans multiple publishing sources. If your task involves mapping citations, tracking funding sources, or analyzing institutional output across the scholarly record, this tool is mandatory. Don't use it if you only need general web search results—it deals exclusively with DOI-registered metadata. Also, don't rely on it to write the actual paper; it

provides the data points (like `get_reference_list` or `search_by_affiliation`), but your agent does the synthesis.

Sifting Through Academic Records is a Full-Time Job

Today, gathering comprehensive academic data feels like detective work. You start by pasting a DOI into one database for citations; then you have to switch to another site to check the journal's reputation using its ISSN. If you need to track who funded the research, you jump to yet a third portal, copying IDs and running multiple searches just to get a basic picture of the field.

With this MCP, your agent handles the complexity. You simply ask for the data point—like 'Give me all works by NIH on protein folding.' The result isn't three separate PDFs; it's one clean list with full metadata, ready for you to use immediately.

Get Full Bibliographic Records with Resolve DOI

The biggest time sink is the manual process of resolving a paper. You might copy an author's name and have no idea if they published under a different alias or institution, forcing you to start over. Checking every reference list manually for completeness is nearly impossible.

Now, using `resolve_doi` instantly structures all that information. It gives you the title, authors, journal, and date in one go. You stop guessing and start analyzing.

Stanford CrossRef MCP: 16 Tools for Scholars

Use these tools to query massive databases of scholarly works, retrieving everything from citation counts to complete bibliographies using structured data requests.

#	TOOL	DESCRIPTION
01	get_citations_count	Quickly determines how many other academic works have cited a specific DOI, helping assess its influence.
02	get_funder_works	Retrieves all published articles linked to a specified funding organization ID, mapping research output back to the source of money.
03	get_journal	Pulls comprehensive metrics on an academic journal by its ISSN, including total article count and quality flags.
04	get_journal_works	Finds articles published within a specific journal, optionally allowing filtering by text search queries.
05	get_publisher	Gathers key metrics on an academic publisher, showing their total output count and metadata coverage score.
06	get_reference_list	Extracts the complete bibliography or list of cited references from a given DOI, including authors and titles.
07	resolve_doi	The primary tool that takes any DOI and returns all structured bibliographic information for that work.
08	search_by_affiliation	Locates academic works by searching using institutional names like 'MIT' or 'Stanford University'.
09	search_by_orcid	Finds a researcher's entire publication history across all journals using their unique ORCID ID.
10	search_funders	Lists and searches for major global funding organizations, such as the National Institutes of Health (NIH).
11	search_journals	Searches for academic journals by title or ISSN to gather metrics and identify journal identifiers.
12	search_preprints	Finds early versions of research papers (preprints) registered across major preprint servers like bioRxiv.
13	search_publishers	Lists and searches for academic publishers, providing their DOI prefixes and total body of work counts.

#	TOOL	DESCRIPTION
14	search_recent_works	Monitors the stream of newly registered DOIs to keep track of the very latest publications across all domains.
15	search_works	Searches through 150M+ academic works using advanced filters, sorting, and text queries.
16	validate_doi	Confirms whether a DOI is active and registered in the system, providing basic metadata if it's valid.

See It in Action

Real prompts you can use once this MCP is connected to your AI agent through Vinkius Cloud.

U Resolve DOI 10.1038/s41586-021-03819-2

I've resolved this DOI via CrossRef. It is: "Highly accurate protein structure prediction with AlphaFold" by Jumper et al., published in Nature (2021). It has been cited over 15,000 times.

U Find works funded by the National Institutes of Health on gene therapy

I've searched for NIH-funded works on gene therapy via CrossRef. Results include recent publications in Nature Medicine, Science, and Cell acknowledging NIH grants.

U Look up all publications by ORCID 0000-0002-8350-519X

I've retrieved all DOI-registered works for this ORCID via CrossRef, showing titles, journals, publication dates, and citation counts for each.

Frequently Asked Questions

01 **How can I find all papers by a specific professor using Stanford CrossRef MCP?**

You use `search_by_orcid`. Providing the researcher's ORCID ID allows the MCP to pull every piece of published work linked to that unique identifier, regardless of which journal printed it.

02 **Does Stanford CrossRef MCP include preprints from non-academic sources?**

No. However, you can use `search_preprints` to find DOIs registered with major preprint servers (like bioRxiv). This ensures you capture cutting-edge research before it enters formal journals.

03 **What is the difference between searching by ISSN and using get_journal?**

Searching by `search_journals` helps you *find* the journal's metrics (like total output). Using `get_journal` then retrieves those specific, detailed quality flags and coverage percentages for a known ISSN.

04 **Can I use Stanford CrossRef MCP to check if a DOI is real?**

Yes. Running `validate_doi` checks the system's registry. It tells you immediately if the identifier exists, and provides basic metadata if it passes validation.

Go Live in 60 Seconds

Get your connection token from **cloud.vinkius.com**, then paste the endpoint URL into any MCP-compatible client.

YOUR MCP ENDPOINT

```
https://edge.vinkius.com/[TOKEN]/mcp
```

CLIENT	WHERE TO CONFIGURE
Claude AI	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
Cursor	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
VS Code	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add `"stanford-crossref": { "url": "…" }`
Windsurf	MCP Settings → `mcp_settings.json` → Add endpoint URL
ChatGPT	Settings → Tools & plugins → Add MCP server → Paste endpoint
Gemini	Extensions → Add MCP Server → Paste endpoint URL

ASK AN AI ABOUT THIS

Let your preferred AI explain this MCP server

- Ask ChatGPT
- Ask Claude
- Ask Perplexity
- Ask Gemini
- Ask Grok

READY TO CONNECT

Stanford CrossRef is live on Vinkius Cloud.

Get your connection token, paste it into your AI agent, and start building. No SDK. No deployment. Just results.

Start at cloud.vinkius.com →

vinkius.com · support@vinkius.com

INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by Stanford CrossRef. All third-party trademarks, logos, and brand names are the property of their respective owners. Their use in this document is strictly for informational purposes to identify service compatibility and interoperability.

DOCUMENT INFORMATION

Generated	June 2026
MCP Server	Stanford CrossRef MCP
Server ID	`019dea60-25c8-739c-ac0a-f2633e810a16`
Platform	Vinkius Cloud for AI Agents
Endpoint	`https://edge.vinkius.com/{token}/mcp`

LICENSE & USAGE

This document is generated automatically by the Vinkius PDF Engine. Content reflects the MCP server configuration at the time of generation and may change as updates are deployed. For the most current information, visit vinkius.com/mcp/stanford-crossref.