MCP SERVER | NO CODE | CLOUD HOSTED

CrossRef MCP for AI Agents

Find and verify scholarly metadata, citations, and academic works by DOI

CrossRef lets your AI client search over 150 million academic works in one place. You can find detailed metadata—like citation counts, authors, publication dates, and funding sources—for journal articles, books, and datasets instantly.

A+ Quality Score 98.33/100

academic-metadata | doi-lookup | scholarly-publishing | citations | research-tools

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

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.

CrossRef MCP

13 tools available
Cloud-hosted on Vinkius

Need to track down scholarly details across massive databases? This MCP connects you directly to CrossRef, giving your AI client access to the world's largest repository of academic metadata. You don't need API keys or to navigate complex web forms; you just ask your agent a question in natural conversation.

Ask it for papers on climate modeling from 2018, and it finds them. Ask it to check which publishers are associated with biomedical journals, and it gives you the list. The result is structured data: titles, authors, DOIs, citation counts, and abstracts—all ready for your next task. If you're building a system that needs reliable, verifiable scholarly information, connecting via Vinkius allows any MCP-compatible client to access this power instantly.

Core Capabilities

01 — Search academic works

Find specific research papers by title, author, keywords, or date.

02 — Look up DOI details

Get complete metadata for any published work using its unique digital identifier (DOI).

03 — Identify journals and publishers

Browse lists of academic journals or find information about specific publishing houses.

04 — Trace research funding

Discover the organizations that funded a particular piece of research, including work counts and locations.

05 — Analyze publication types

Filter searches to focus only on specific kinds of academic output, like book chapters or datasets.

One Click on Vinkius — From Prompt to Execution

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

01 Connect your preferred AI client via Vinkius and enable this MCP.

02 Instruct your agent with a query, such as 'Find all papers on quantum computing from 2023 funded by the DOE.'

03 The MCP processes the request and returns structured data containing titles, authors, DOIs, and full metadata.

The bottom line is, you stop doing manual searches across different academic sites; your agent does it for you.

Built For

This MCP changes the day-to-day life of anyone working with published research. If you're a researcher drowning in citation checks, or a librarian needing to verify journal details quickly, this is built for your pain points.

Academic Researcher

Needs to find related papers and track down the funding sources behind a specific piece of research.

University Librarian

Requires quick verification of DOIs, checking journal ISSNs, or mapping out publisher relationships for resource acquisition.

Scientific Writer

Must gather comprehensive background information, pulling together diverse metadata points like work types and license details for a manuscript.

What Changes When You Connect

01 Instantly get complete data on any paper's origin. Instead of guessing, use the `get_work` tool to pull authors, abstract, citation count, and journal info from a single DOI.

02 Stop manually cross-referencing journals. Use `search_journals` to find ISSNs and work counts for any academic publication you are researching.

03 Track down research funding sources accurately. If you know the topic but not the money trail, use `search_funders` to identify key supporting organizations.

04 Build specialized filters into your workflow. Use `get_types` to ensure you only search for specific deliverables, like datasets or book chapters, filtering out noise.

05 Batch process metadata retrieval. When reviewing many related papers, pass them all at once using `get_works_by_doi_list` instead of running individual lookups.

Real-World Applications

Verifying a bibliography for publication

A scientific writer needs to confirm every citation in a manuscript. They ask their agent to use `get_work` repeatedly on all DOIs, ensuring every paper has the correct author list and current citation count before submitting.

Mapping academic relationships

A librarian needs to understand which publishers are associated with specific journals. They run a search using `search_publishers` combined with `search_journals` to map the entire ecosystem of knowledge resources at an institution.

Investigating research impact

A researcher finds a key paper and needs context. They ask their agent to use `get_funder` on the DOI's metadata, revealing which organizations initially financed that groundbreaking work.

Broad topic literature review

A student is writing a thesis on AI ethics. Instead of manual searches, they use `search_works` with free-text queries and advanced filters (like publication date or author) to pull a broad set of relevant papers for their bibliography.

Patterns to Avoid

Trying to search by generic terms

✗ AVOID

Asking the agent, 'Show me all research about climate.' This returns too much noise and lacks structured data.

✓ INSTEAD

Use `search_works` with specific filters. Instead of a broad query, narrow it: 'Search for works titled X, published in 2022, by author Y' to get actionable metadata.

Assuming all data is available

✗ AVOID

Believing that every piece of research has its funding source listed. This leads to incomplete analysis.

✓ INSTEAD

Always use `get_funder` or `search_funders` after finding a work's DOI. These tools confirm known financial support for the specific paper.

Treating it like a simple search engine

✗ AVOID

Using the MCP just to find links, ignoring structured data points like ISSN or citation counts.

✓ INSTEAD

Always follow up your initial query with entity-specific tools. After finding a journal, run `get_journal` for its full details; don't rely only on the search result summary.

The Right Fit

Use this MCP if your research requires verifiable academic data points: specific citation counts, ISSN numbers, or funding bodies. If you need to know *who* wrote it and *where* it came from, use CrossRef. Don't use this if you just need general web content—it's for structured scholarly metadata. If your goal is simply finding a PDF link, crossref isn't the main tool; but if you need the full context (authors, journal name, work count), then running `get_work` or `search_works` gives you that depth. Remember, this MCP excels at linking disparate pieces of scholarly information together.

CrossRef: Mastering Scholarly Metadata Retrieval with CrossRef and AI Agents

The current process for academic research is a mess of tabs and manual copy-pasting. You find one great paper, then you have to jump to three different websites—one for the journal details, one for the funding source, and another just to check the citation count. This takes hours and introduces errors.

With this MCP, your agent handles all that jumping. Find a DOI, feed it into your client, and immediately get back structured metadata: who published it, what journals are involved, and how many times it's been cited. You get verified data points in one go.

CrossRef: Verifying Journal Details and Publication Types with CrossRef MCP

Before submitting a paper or starting a literature review, you're usually forced to manually look up journal ISSNs or determine if the work is a book chapter versus a full article. This verification process slows down every step.

Now, your agent uses tools like `search_journals` and `get_types`. It instantly provides the official ISSN for any journal name you give it, or confirms whether the source material was originally published as a 'proceedings' paper. You skip the guesswork entirely.

CrossRef: 13 Tools for Scholarly Metadata Search and Citation Lookup

Use these tools to perform highly specific searches across academic journals, funders, publishers, and individual research works.

#	TOOL	DESCRIPTION
01	get_funder	Retrieves specific details about a research funding organization, including their location and total works funded.
02	get_journal	Gets comprehensive data for an academic journal using its ISSN, including publisher and work count.
03	get_licenses	Provides a list of all known academic licenses used in scholarly works, along with their names and IDs.
04	get_member	Retrieves details for a specific CrossRef member organization, including its prefixes and total work count.
05	get_publisher	Gets information on an academic publisher, showing their name, work count, journal list, and location.
06	get_types	Provides a full list of recognized publication types, useful for narrowing down search results.
07	get_work	Pulls all metadata—title, authors, abstract, citation count, and journal details—for a single academic work given its DOI.
08	get_works_by_doi_list	Gets metadata for multiple works at once by providing a comma-separated list of DOIs.
09	search_funders	Searches across research funders to find names, IDs, and how many works they have supported.
10	search_journals	Looks up academic journals by name to get their ISSNs, publisher details, and total work count.
11	search_members	Searches the catalog of CrossRef member organizations by name or ID.
12	search_publishers	Finds academic publishers by name to get their work count, journal list, and organization details.

#	TOOL	DESCRIPTION
13	search_works	Performs comprehensive searches across the entire database using keywords, author names, or specific DOIs.

See It in Action

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

U Find papers on large language models from 2023 and show me their citation counts.

Search Results: LLMs (2023)

TITLE	AUTHORS	JOURNAL	YEAR	CITATIONS
The Transformer Architecture Revisited	Vaswani et al.	Nature Machine Intelligence	2023	12,450+
Context Window Limits	Smith & Jones	IEEE Trans. AI	2023	8,900+

Note: Citation counts are dynamic and retrieved live.

U What kind of funding supported the work by DeepMind on protein folding?

Funding Details for AlphaFold:

This research was primarily funded by:

- **Organization Name:** National Institutes of Health (NIH)
- **Country:** USA
- **Total Works Funded:** 312
- **Location:** Bethesda, MD

The NIH is a key contributor to this field.

U List the biggest academic publishers and how many journals they manage.

Top Academic Publishers by Scope:

1. **Elsevier:** 450+ Journals (Focus: Biomedicine)
2. **Springer Nature:** 380+ Journals (Focus: STEM & Humanities)
3. **Wiley Online Library:** 190+ Journals (Focus: Social Sciences)

These counts reflect the total number of unique journal titles they manage.

Frequently Asked Questions

01 **How can CrossRef help me verify citation details for my bibliography?**

CrossRef provides accurate, verifiable metadata for every academic work. By using tools like `get_work`, you get the exact authors, journal name, and up-to-date citation counts, ensuring your bibliography is flawless.

02 **Do I need to know a DOI to use CrossRef with my AI agent?**

No. While providing a DOI gives you the most precise result, this MCP allows you to search using natural language queries based on keywords, authors, or publication dates.

03 **Can I find out who paid for a specific piece of research using CrossRef?**

Yes. The MCP includes tools to trace the funding sources. You can discover which organizations funded the work and how many other projects they've supported, giving you full context.

04 **Is this better than just searching Google Scholar for papers?**

This is more powerful because it delivers structured data directly to your agent. Instead of a link list, you get specific fields like ISSNs and work counts that are ready for immediate use in reports.

05 **How do I find out which types of publications exist (like datasets)?**

You can run a simple query using the `get_types` tool. It gives you a clean, comprehensive list of all recognized publication formats, helping you filter your searches precisely.

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 `"crossref-alternative": { "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

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 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	CrossRef MCP
Server ID	`019d842b-485f-7082-a2ee-0520de240e4c`
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/crossref-alternative.