MCP SERVER | NO CODE | CLOUD HOSTED

CrossRef MCP for AI Agents

Access the World's Largest Registry of Scholarly Metadata

CrossRef gives your AI agent direct access to the world’s largest registry of scholarly metadata, covering over 140 million records across every scientific discipline. Instantly resolve DOIs, track citation counts, or search for all publications from a specific researcher—all in one place.

A+ Quality Score 98.33/100

scholarly-metadata | doi-resolution | academic-research | citation-tracking | bibliographic-data | search-api

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

CrossRef MCP

3 tools available
Cloud-hosted on Vinkius

This MCP connects your AI agent directly to CrossRef, giving it visibility into the world's most comprehensive database of scholarly metadata. You stop hunting across multiple academic indexes and start asking questions that get definitive answers.

Whether you're writing a literature review or building an analytical model, you can use this MCP for immediate data retrieval. Your agent can find any published work—journals, books, datasets, conference papers—using free-text queries to map out entire fields of research. If you just have a DOI, the tool immediately pulls all the deep metadata: title, full author list, journal details, and how many times it has been cited.

It's built for people who need verifiable data fast. Because this MCP is hosted on Vinkius, your agent connects once and gains access to cross-disciplinary research tools without needing any specific API keys or manual setup steps. It lets you verify citations and build bibliographies against the definitive source of academic publishing.

Core Capabilities

01 — Search scholarly works by topic

Your agent can search across 140M+ records for any published work, including DOI and citation metrics.

02 — Resolve metadata using a DOI

The MCP pulls complete details—authors, journal, year, type, and citation count—for any given Digital Object Identifier (DOI).

03 — Gather all publications by an author

Your agent finds every published work associated with a specific researcher across major publishers.

One Click on Vinkius — From Prompt to Execution

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

01 Connect your AI client to this MCP in Vinkius.

02 Ask your agent to perform an action, like finding all works by a named author or resolving a specific DOI number.

03 The agent executes the request and returns structured data containing full metadata, citation counts, and bibliographic details.

The bottom line is that you pass scholarly reference points (like DOIs or names) to your agent, and it pulls back all the associated academic context immediately.

Built For

This MCP is essential for anyone who works with published knowledge. It's for data analysts needing citation metrics, science writers verifying facts, or academics building comprehensive bibliographies without leaving their AI client.

Academic Researcher

Uses the tool to verify citations and build detailed literature reviews by looking up specific DOIs or tracing an author's entire publication history.

Science Writer / Journalist

Retrieves instant, verifiable details about a scientific paper—like its original journal and citation count—to ensure factual accuracy in articles.

Data Analyst (Bibliometrics)

Explores the academic landscape by running broad searches with metrics to analyze trends or identify key contributors in a field.

What Changes When You Connect

01 Verify facts instantly. Instead of guessing, use `get_crossref_doi` to pull precise details—like journal name and year—for any research paper with a DOI.

02 Build comprehensive bibliographies faster. The MCP lets you track down every publication associated with an author by running the `search_crossref_author` tool, giving you a complete history of work.

03 Map out entire fields of study. Use `search_crossref` to run broad queries across 140M+ records and gather citation metrics for bibliometric analysis in one go.

04 Consolidate data sources. You don't need separate tools for books, journals, or datasets; this MCP handles all scholarly types simultaneously.

05 Save time on manual lookup. The agent resolves the complexity of academic publishing so you get structured data right away.

Real-World Applications

Checking a source in a news piece

A science journalist needs to verify a claim about CRISPR technology. They ask their agent to use `get_crossref_doi` on the provided DOI, and it immediately confirms the paper's original journal, full author list, and how many times the finding has been cited.

Writing a literature review for a thesis

A PhD candidate is writing about deep learning. They ask their agent to use `search_crossref` with 'deep learning' as the query, getting hundreds of results that include citation counts and full bibliographic data, helping them prioritize key foundational papers.

Tracking a researcher's impact

A principal investigator needs to know all work by a collaborator. They use `search_crossref_author` with the name, receiving a ranked list of every publication, sorted by relevance and total citations.

Building an internal knowledge graph

A data analyst wants to map relationships between concepts in quantum computing. They use `search_crossref` repeatedly on different keywords, gathering structured metadata and citation links for a robust internal resource.

Patterns to Avoid

Relying only on general search engines

✗ AVOID

Pasting a handful of paper titles into Google Scholar and manually copying the authors, years, and citations into a spreadsheet. This process is slow and often misses key metadata.

✓ INSTEAD

Instead, give your agent the DOI or author name and let it use `get_crossref_doi` or `search_crossref_author`. You get all the reliable data points in one structured output.

Assuming citation completeness

✗ AVOID

Thinking that a search result is definitive because it only lists the title and authors. Often, crucial details like the publication type or total citation count are missing.

✓ INSTEAD

Always use `search_crossref` to query across the entire 140M+ registry. Every result includes DOI resolution and the full bibliographic data you need.

Confusing general search with metadata lookup

✗ AVOID

Searching for 'COVID-19' but not specifying a piece of research, resulting in thousands of irrelevant links. You waste time sifting through non-academic web pages.

✓ INSTEAD

If you know the DOI or author name, use `get_crossref_doi` or `search_crossref_author`. These tools pinpoint exact academic records against the definitive registry.

The Right Fit

Use this MCP if your primary need is verifiable scholarly metadata. If you have a specific DOI, author name, or general research topic and need full citation metrics, history, and comprehensive details—this is it. Don't use it if you just need to find a simple website or check business hours; those are outside its scope. You shouldn't use this MCP if all you want is an abstract summary without source verification, because the goal here is always deep, verifiable data retrieval. If your task is solely content generation from scratch (like writing a blog post based on common knowledge), then other general LLMs suffice. But when accuracy and attribution are non-negotiable, CrossRef provides the necessary rigor.

CrossRef MCP for AI Agents: Resolving Citation Data in Academic Research

Today, verifying academic facts is a huge pain point. Researchers spend hours copy-pasting DOIs or titles into different databases just to get the full author list and citation count. They jump between Google Scholar, university repositories, and publisher sites—it's a massive time sink that requires tedious manual checking.

With this MCP, your agent handles all of that legwork for you. By running `get_crossref_doi` or using general searches, you get the definitive source data instantly. You don't just get metadata; you get certainty.

CrossRef MCP for AI Agents: Tracing Author Output in Scholarly Metadata

If you’re trying to build a bibliography or map an author's contribution, the manual process is awful. You have to search by name repeatedly across different platforms and manually track which works are related.

This MCP solves that using `search_crossref_author`. It pulls every work associated with one person into a single, comprehensive list, ranked by relevance and providing reliable citation metrics for everything.

3 CrossRef Tools for Academic Bibliographic Data Retrieval

Use these tools to search across 140M+ records, look up specific DOIs, or trace an author's full publication history.

#	TOOL	DESCRIPTION
01	search_crossref_author	Finds all publications by a given author name and returns them sorted by relevance along with their total citation counts.
02	get_crossref_doi	Looks up complete metadata (title, authors, journal, year, citation count) instantly using a specific DOI number.
03	search_crossref	Searches 140M+ scholarly works across all scientific disciplines, returning DOI and full bibliographic data for any topic.

See It in Action

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

U What are the key papers on mRNA vaccine technology?

Key Papers in mRNA Technology

- **Title:** Lipid nanoparticle delivery of mRNA encoding Spike protein
- **Journal:** Nature Medicine (2021)
- **Authors:** Smith J., Lee K., Chen P.
- **Cited:** 987 times
- **DOI:** [Example DOI]

This paper is a foundational piece. It shows the initial work on delivery systems, which was critical for modern vaccine development.

U Find all papers by Jennifer Doudna related to CRISPR.

Publications by Jennifer Doudna (Top 5)

YEAR	TITLE SNIPPET	JOURNAL	CITATIONS
2012	Programmable Dual-RNA-Guided...	Science	15,000+
2017	CRISPR/Cas9 in vivo applications	Nature	4,500+
2020	Genome editing for therapeutics	Cell	1,200+

The results confirm her primary contributions and show the trajectory of research over time.

U Give me all the details on this DOI: 10.1038/nature12373

Metadata Lookup Complete

- **Full Title:** Sequence-specific antimicrobials using efficiently delivered RNA-guided nucleases
- **Authors:** David Bikard, et al.
- **Journal:** Nature Biotechnology
- **Published:** 2014
- **Type:** Journal Article
- **Total Citations:** 682

The record is solid. The full text link resolves to the publisher page.

Frequently Asked Questions

01 **How does CrossRef MCP help me find reliable academic sources?**

It gives you access to a massive, definitive registry of scholarly metadata. Instead of relying on general web searches, the MCP connects your agent to structured data, providing verifiable DOIs and full citation counts for every source.

02 **I need to track an author's career history; can CrossRef MCP do that?**

Yes. You can use the MCP to find all publications by a specific researcher across major publishers simultaneously. It generates a comprehensive, ranked list of their work and citation impact.

03 **What if I only have a DOI number? Can CrossRef MCP still give me enough detail?**

Absolutely. If you provide just the DOI, the MCP pulls out every piece of associated metadata: the full title, all authors, the journal it appeared in, and its total citation count.

04 **Is CrossRef MCP better than Google Scholar for bibliography data?**

Yes. While Google Scholar is good for discovery, the MCP accesses a professional registry with definitive metadata resolution. It provides structured data points needed for reliable bibliometric analysis and academic writing.

05 **Does CrossRef MCP cover datasets or just journal articles?**

No, it's universal. The system searches across journals, books, conference papers, *and* datasets. It is designed to map the entire scholarly landscape, not just published texts.

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": { "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	`019d757f-02a7-7015-bd73-f35632d2ee53`
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.