MCP SERVER | NO CODE | CLOUD HOSTED

OpenAlex MCP

Map the global research landscape with natural conversation.

OpenAlex connects your AI agent directly to a massive, open index of global scholarly research. It lets you search millions of works, authors, institutions, and funding sources using natural conversation. Stop jumping between databases; find metadata on anything from academic papers to university affiliations instantly.

A+ Quality Score 100/100

scholarly-data | academic-research | bibliometrics | metadata-search | open-science | research-automation

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

OpenAlex MCP

14 tools available
Cloud-hosted on Vinkius

Need to map out the full landscape of a specific scientific field? OpenAlex gives your AI agent access to an open index of global research data, tracking everything published across academia. You can ask for works related to a topic and get back structured metadata, or you can identify key researchers by name and pull their entire academic profile, including all major affiliations. This MCP lets you explore the connections between institutions, the people who work there, and the money that funds it. Whether you're tracking citation trends across different disciplines or just verifying an author's credentials, your agent handles the heavy lifting. Connect this MCP through Vinkius to give any compatible client access to this deep knowledge base.

Core Capabilities

01 — Search and filter academic papers

You can list, search, or filter scholarly works, including articles, books, and datasets based on criteria like author ID, publication year, or topic.

02 — Identify researcher profiles

Retrieve detailed information about any academic author, showing their primary affiliations and overall impact metrics.

03 — Map out institutional connections

List or fetch details on universities and research organizations globally, helping you understand the geographic spread of research.

04 — Trace funding sources

Find single funder profiles or list multiple funders associated with specific research projects.

05 — Explore source details

Get metadata for journals, conferences, and repositories (sources) or discover the core topics driving a field of study.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP on Vinkius. You can optionally enter your OpenAlex API Key if you need faster rate limits.

02 Tell your AI client what specific academic data you need, like 'Find all papers about quantum computing from MIT in 2021.'

03 The agent runs the necessary search through this MCP and returns structured metadata for the works, authors, or institutions you asked about.

The bottom line is your AI client reads complex academic indexes using natural language queries, returning clean data points instantly.

Built For

This MCP is for academics and analysts who spend too much time manually cross-referencing citation databases. Use this if you need to map the provenance of an idea or track a researcher's career trajectory across multiple publications.

Academic Researcher

Uses list_works and get_author to quickly find relevant literature, verify citations, and understand which key authors published specific papers.

Data Scientist (Bibliometrics)

Runs complex queries using list_institutions or list_funders to gather data on research trends, map academic collaboration networks, and analyze funding patterns programmatically.

Research Analyst

Uses get_source and get_topic to trace the history of a subject by identifying core journals, major contributing sources, or emerging fields of study.

What Changes When You Connect

01 Stop copy-pasting author names. Use list_authors first to resolve a researcher's name into a verifiable ID, then use get_author to pull their complete profile in one step.

02 Track funding history easily. If you know the work, you can check which funders were involved using get_funder; if you know the funder, you can list what they supported via list_funders.

03 Build research maps instantly. Instead of visiting separate university directories, use list_institutions to find top academic hubs and then query them for related works using list_works.

04 Pinpoint source quality. Need to know if a paper is from a major journal or an obscure repository? Use get_source to check the full details of the publication venue.

05 Navigate complex topics without specialized knowledge. Use list_topics and search by category to understand the foundational subjects driving current academic research.

Real-World Applications

Mapping a field’s origins

A bio-tech analyst needs to know which institutions first studied CRISPR gene editing. They ask their agent, 'List works about CRISPR and filter by institution.' The agent uses list_works combined with list_institutions to return the top five universities and the foundational papers published there.

Verifying a cited paper’s context

A student is writing a thesis and finds an author mentioned. They ask their agent, 'What are this author's key affiliations and what sources did they publish in?' The agent uses get_author to confirm the researcher’s profile and list_sources for primary publication venues.

Tracing academic influence

A data scientist wants to see which organizations are leading research on clean energy. They prompt, 'List institutions that published works related to solar power over the last decade.' The agent uses list_works and list_institutions to return a prioritized list of contributors.

Analyzing grant dependencies

A consultant needs to know if a recent paper was funded by private or governmental sources. They ask, 'For this work ID, who provided the funding?' The agent executes get_work and then uses get_funder to provide the specific corporate or government entity responsible for the research.

Patterns to Avoid

Searching by vague keywords

✗ AVOID

Just asking 'Tell me about quantum physics.' This gets a huge, unmanageable list of general results and doesn't pinpoint specific works or authors.

✓ INSTEAD

Be precise. Use the tool to narrow your search: 'List works on quantum computing published between 2018 and 2023 by institutions in Germany.' This uses list_works with filtering criteria.

Ignoring ID lookups

✗ AVOID

Trying to query an author's work using only their name repeatedly. The AI client can get confused or pull too many irrelevant results.

✓ INSTEAD

Always resolve first. Use list_authors to convert the name into a stable ID, then use that ID with get_author or list_works for accurate, targeted data.

Assuming single-tool function

✗ AVOID

Asking only 'List all works.' This returns millions of records and requires manual filtering on the user's end.

✓ INSTEAD

Always refine your query. Use list_works with specific filters like `publication_year:>2022` or `author.id:A123` to get a manageable, actionable dataset.

The Right Fit

Use this MCP if your research requires deep, structured metadata retrieval across the entire academic ecosystem. You need to map connections—who funded what, who collaborated with whom, and which source is most authoritative for a given topic. Don't use it if you are looking for general news or casual facts; OpenAlex only handles peer-reviewed, indexed scholarly data. If your goal is simply to search Wikipedia or find public domain creative works, this MCP

isn't the right fit. However, if you need to validate an academic claim—for example, confirming a paper's existence (get_work), checking its source (get_source), and identifying its primary contributor (get_author)—this is exactly what you need.

Keeping track of the full history of research sources is a nightmare.

Right now, if you want to know where a key academic paper came from, you have to jump through multiple portals: checking databases for authors, going to university websites for affiliations, and then searching journal sites just to see the source. You copy-paste identifiers dozens of times until you piece together who funded it, which institution published it, and if the data is even complete.

With this MCP, your agent handles all those hops. You tell it what you need—say, 'What sources are key for AI ethics?' The system runs complex queries across list_sources, get_topic, and list_institutions to build a comprehensive picture of the field's intellectual origins.

The OpenAlex MCP delivers complete academic profiles with get_author.

Manual profiling means opening one tab for affiliations, another for citation counts, and a third just to list the works. You lose hours cross-referencing these details, and you might miss key publications from secondary sources.

Using this MCP, your agent pulls everything together automatically. It resolves an author's name into their ID and executes get_author, giving you a single, clean summary of their entire career history without leaving the chat.

OpenAlex: 14 Tools for Academic Research

Use these tools to query the OpenAlex database directly. Search academic papers, profiles, organizations, and more with specific functions.

#	TOOL	DESCRIPTION
01	`get_author`	Retrieves all profile data for one specific author using their OpenAlex ID.
02	`get_funder`	Gets the full profile of a single funder organization by its unique OpenAlex ID.
03	`get_institution`	Pulls all available details for one university or research body using its OpenAlex ID.
04	`get_publisher`	Retrieves the profile information for a single publishing house by its OpenAlex ID.
05	`get_source`	Gets metadata about one specific journal, conference, or repository using its OpenAlex ID.
06	`get_topic`	Retrieves detailed information for a single defined academic topic area by its OpenAlex ID.
07	`get_work`	Fetches the full metadata record for one scholarly work, such as an article or dataset, using its unique ID.
08	`list_authors`	Searches and filters authors by name, affiliation, or other criteria to resolve names into IDs for later use.
09	`list_funders`	Lists, searches, and applies filters across multiple funding organizations.
10	`list_institutions`	Provides a searchable directory of academic institutions and research groups worldwide.
11	`list_publishers`	Searches and filters publishing entities by name or scope.
12	`list_sources`	Lists, searches, and filters academic sources like journals or conference proceedings.
13	`list_topics`	Provides a searchable list of recognized research topics and subject areas.
14	`list_works`	Searches, filters, sorts, and groups large sets of scholarly works (articles, books, datasets) by multiple criteria.

See It in Action

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

U Search for scholarly works about 'Large Language Models' published in 2023.

I've found several highly cited works from 2023. Top results include 'Sparks of Artificial General Intelligence' and 'LLaMA: Open and Efficient Foundation Language Models'. Would you like the full metadata for any of these?

U List the top 5 research institutions in Germany by works count.

Based on OpenAlex data, the top institutions in Germany are: 1. Helmholtz Association, 2. Max Planck Society, 3. University of Munich, 4. Heidelberg University, 5. Technical University of Munich. Would you like to explore the authors associated with any of these?

U Get detailed information for author ID A5023888321.

Retrieving profile... This ID belongs to Yoshua Bengio. He has over 1,000 works indexed, with a high citation count in Computer Science and Artificial Intelligence. His primary affiliation is Université de Montréal.

Frequently Asked Questions

01 **How do I use OpenAlex MCP to find works by a specific university?**

You can list institutions first using list_institutions to get the ID for that university. Then, run list_works and filter results by including that institution's ID in your search query.

02 **Can OpenAlex MCP help me track funding sources?**

Yes. You can use list_funders to search for organizations or get details on a specific funder using get_funder. This helps you map the financial support behind academic research.

03 **What's the difference between list_works and get_work in OpenAlex MCP?**

list_works is for searching large sets of papers based on criteria like year or topic. get_work requires you to already have a specific work ID and retrieves its full metadata record.

04 **Does the OpenAlex MCP support filtering by publication date?**

Yes, list_works supports advanced filtering, including setting minimum (`:>`) or maximum (`:<`) years for the publication year. This is critical for tracking research trends over time.

05 **Do I need to know the OpenAlex ID to use get_author?**

While it's most accurate, you can first use list_authors to search by name or affiliation and retrieve the necessary author ID before running get_author.

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 `"openalex-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

OpenAlex 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 OpenAlex. 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	OpenAlex MCP
Server ID	`019e38cd-4711-72f7-999e-65c2fd419828`
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/openalex-alternative.