MCP SERVER | NO CODE | CLOUD HOSTED

ScienceDirect MCP

Access millions of scientific articles instantly.

Elsevier ScienceDirect connects your AI agent directly to millions of scientific, technical, and medical articles and book chapters. Search for specific data using advanced queries, pull full text or abstracts instantly, and verify access rights—all from one chat window.

F Quality Score 7.21/100

scholarly-articles | peer-reviewed | academic-search | research-data | metadata-retrieval | scientific-literature

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

SCAN TO VIEW

The infrastructure that powers AI agents in the real world.

Vinkius connects AI to the world's software through secure, enterprise-grade infrastructure — enabling real-world execution at scale, built on the Model Context Protocol (MCP).

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the cloud infrastructure where AI agents connect 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.

Elsevier ScienceDirect MCP

9 tools available
Cloud-hosted on Vinkius

When you connect this MCP, your AI client becomes a specialized research assistant trained on the world's foremost academic literature. You stop clicking through library portals and start asking questions. Instead of wading through paywalls or navigating complex databases manually, you ask for specific information—say, 'What are the key findings regarding CRISPR applications in human tissue?' The agent finds the articles, pulls the relevant abstracts, and hands them back to you, ready for analysis. Furthermore, it helps manage institutional compliance by verifying whether you have the right to host or use a particular paper. This deep connection allows your AI client to perform complex research tasks that used to require multiple specialized tools working together. By hosting this MCP on Vinkius, you gain access to a single point of entry for global scientific data.

Core Capabilities

01 — Search the entire database

You can execute advanced searches across articles and books using Boolean logic to narrow down results quickly.

02 — Fetch article content

The agent pulls either a full text or a concise abstract of an academic paper based on identifying details like DOI or Scopus ID.

03 — Check data and hosting rights

You can verify if your institution has permission to host content or check for active embargo periods on specific articles.

04 — Identify source types

The agent fetches specialized metadata for both journal titles (by ISSN) and books (by ISBN).

One Click on Vinkius — From Prompt to Execution

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

01 First, subscribe to this MCP in the Vinkius catalog and provide your Elsevier API Key along with any institutional tokens.

02 Next, tell your AI client exactly what you need—for example, 'Search for articles on deep learning in oncology.'

03 Finally, the agent processes the query against ScienceDirect, returning structured data like abstracts, metadata, or full text directly to your chat environment.

The bottom line is that your AI client bypasses web portals and speaks directly to the scientific database's core systems.

Built For

This MCP is essential for anyone whose job requires deep, rapid access to scholarly knowledge. It's for researchers who are tired of spending hours manually tracking down abstracts and verifying citation details across multiple paywalled sites.

Academic Researcher

They use this MCP to automate literature reviews, gathering full texts or metadata on dozens of papers in minutes instead of days.

Librarian/Archivist

They rely on the ability to verify hosting permissions and check article entitlements for institutional compliance when curating digital collections.

Data Scientist

They use it to systematically collect structured metadata, like subject classifications or citation details, to build comprehensive literature datasets.

What Changes When You Connect

01 Instead of manually copying identifiers, you can use the `get_article` tool to pull abstracts or full texts using just a DOI or Scopus ID. Your agent handles the heavy lifting so you don't have to.

02 You eliminate compliance risk by running checks like `get_hosting_permissions`. This ensures that any content you plan to use, cite, or archive is legally available for your institution.

03 The platform gives you structured data immediately. Use `get_article_metadata` to pull detailed background info on a paper, giving you fields like subject classifications and author details in an easy-to-use format.

04 You can manage the entire scope of knowledge. You don't just search articles; use `get_serial_title` for journals (by ISSN) or `get_nonserial_title` for books (by ISBN) to map out your source material efficiently.

05 The system supports complex querying. The `search_sciencedirect` tool lets you write advanced Boolean searches across the full text of millions of articles, finding connections you'd miss otherwise.

Real-World Applications

Reviewing a literature gap for a thesis

A PhD candidate needs to prove that nobody has studied topic X in the last five years. They ask their agent to run an advanced search using `search_sciencedirect` with precise date and keyword filters, getting a clean list of results they can analyze immediately.

Curating a departmental digital archive

A librarian needs to ensure all collected papers are reusable. They use the agent to check every title using `get_article_entitlement` and `get_hosting_permissions`, building an instant compliance report for their department.

Analyzing source material types

A data science team is compiling a dataset of sources. They use the agent to check both journal titles via `get_serial_title` and book details via `get_nonserial_title`, ensuring they categorize every piece of source material correctly.

Quickly summarizing complex research

A user receives a DOI for an article they haven't read. They prompt their agent to execute `get_article` using the DOI, instantly pulling a summary or full text so they can decide if it's relevant without spending minutes navigating journal paywalls.

Patterns to Avoid

Over-relying on Google Scholar

✗ AVOID

The user assumes that because the paper is indexed on Google, they can access its full text or metadata anywhere. They get a list of links but no usable data.

✓ INSTEAD

To guarantee you get the clean, structured metadata and check actual institutional access rights, always use `get_article_metadata` or `get_article_entitlement` through this MCP.

Confusing source types

✗ AVOID

The user tries to search for a book title using journal article parameters, leading to vague or incorrect results.

✓ INSTEAD

Be specific. If it's a book, use `get_nonserial_title` with the ISBN. If it's a journal, use `get_serial_title` with the ISSN.

Forgetting data compliance

✗ AVOID

A researcher downloads an article and plans to host it on their university server without checking usage rights, risking legal issues.

✓ INSTEAD

Always run `get_hosting_permissions` before archiving or publicizing material. This verifies the embargo status and hosting requirements for safe use.

The Right Fit

Use this MCP if your research workflow is defined by source verification, deep metadata extraction, or cross-source comparison (e.g., comparing a journal article to a chapter in an ISBN book). You absolutely need it when you must prove the provenance and legal availability of scholarly content. Don't use this if you just need general web search results; for that, a standard web retrieval tool is

fine. However, if your goal is academic rigor—you need structured data like subject classifications via `get_subject_classifications` or guaranteed full-text access to guarantee accuracy—then this MCP is non-negotiable.

Academic research requires a painful amount of manual verification.

Today, if you need to compile a literature review, the process is tedious: You find a promising paper via Google Scholar. Next, you open three browser tabs—the journal site, JSTOR, and your library proxy portal. You click through paywalls until you find the PDF. Then, you copy-paste the DOI into a separate spreadsheet to manually record the subject classifications, the full title of the serial (ISSN), and whether or not it's cleared for public hosting.

With this MCP, that whole multi-tab, copy-pasting nightmare disappears. You ask your agent to find the paper, and in return, you get structured data: The article's full text, its metadata, and a clear status report on whether you can host it—all without opening a single new browser tab.

Get Article Text & Metadata with Elsevier ScienceDirect MCP

You no longer have to jump between different database types. Need details on a journal? Use `get_serial_title`. Need a book's background? Run `get_nonserial_title`. If you just need the raw text, use `get_article` with the correct ID.

The result is reliable data flow. You get immediate access to structured scientific knowledge that previously required multiple logins and specialized portal navigation. It's instant.

Elsevier ScienceDirect: 9 Tools for Research Data

These tools allow your AI client to perform highly specific actions on the ScienceDirect database, handling everything from running advanced searches to verifying article access rights.

#	TOOL	DESCRIPTION
01	`get_article_entitlement`	Checks if a specific user has the rights to view an article.
02	`get_article_metadata`	Retrieves detailed background information about an academic paper.
03	`get_article`	Pulls the full text or just a summary of an article using its ID, DOI, or Scopus number.
04	`get_hosting_permissions`	Determines if an article is currently subject to an embargo or can be hosted in other repositories.
05	`get_nonserial_title`	Gets the title and metadata for books identified by their ISBN number.
06	`get_object`	Finds attached media files or supplemental materials linked to an article.
07	`get_serial_title`	Retrieves the title and metadata for academic journals identified by their ISSN number.
08	`get_subject_classifications`	Fetches the specific subject categories or taxonomy tags assigned to a paper.
09	`search_sciencedirect`	Runs advanced, full-text searches across ScienceDirect's entire publication library.

See It in Action

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

U Search ScienceDirect for recent articles about CRISPR in Title/Abstract/Keywords.

I've found several recent articles on CRISPR. The top results include 'CRISPR-Cas9 Gene Editing in Human Cells' and 'Advances in CRISPR Technology'. Would you like the abstract for any of these?

U Get the full text for the article with DOI 10.1016/j.cell.2023.01.001.

Retrieving full content for DOI 10.1016/j.cell.2023.01.001... I have successfully fetched the article. It covers new insights into cellular signaling pathways. How would you like me to summarize it?

U Check the hosting permissions for PII S0022-2836(21)00123-4.

Checking hosting permissions for PII S0022-2836(21)00123-4... This article has an active embargo period ending on 2024-12-01. After this date, it can be hosted in institutional repositories.

Frequently Asked Questions

01 **How do I use the Elsevier ScienceDirect MCP for full-text articles?**

Use the `get_article` tool. You'll need to provide an identifier like a DOI, Scopus ID, or PII number to pull either the complete text or just the abstract.

02 **Can I use Elsevier ScienceDirect MCP to check book information?**

Yes, you can. Use the `get_nonserial_title` tool and provide the ISBN number; it will retrieve the title and metadata for the book source.

03 **What is the best way to search across many topics with ScienceDirect MCP?**

Use the `search_sciencedirect` tool. It handles advanced Boolean queries, letting you filter by keywords, dates, and specific fields across millions of articles.

04 **Does Elsevier ScienceDirect MCP check if I can host research papers?**

Absolutely. The `get_hosting_permissions` tool checks the article's embargo status and determines if it is cleared for use in institutional repositories, protecting you from compliance issues.

05 **How does ScienceDirect MCP help with subject classification?**

Use the `get_subject_classifications` tool. It automatically pulls the assigned taxonomy tags and subject categories, which is perfect for building structured data sets in your research.

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 `"elsevier-sciencedirect": { "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

Elsevier ScienceDirect 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 Elsevier ScienceDirect. 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	Elsevier ScienceDirect MCP
Server ID	`019e3890-3f29-704a-aaf0-c290f42429bf`
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/elsevier-sciencedirect.