MCP SERVER | NO CODE | CLOUD HOSTED

PubMed Central MCP

Analyze deep biomedical literature and citations.

PubMed Central MCP connects your AI agent directly to PMC, the world's open-access library of biomedical and life sciences literature. Search millions of full-text articles, analyze citation networks, and pull structured data in JSON format without ever leaving your client. This is deep research retrieval for modern science.

B Quality Score 87.3/100

pubmed | biomedical | open-access | academic-search | full-text

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

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.

PubMed Central MCP

7 tools available
Cloud-hosted on Vinkius

When you need to understand a complex scientific topic, you shouldn't have to manually download PDFs and cross-reference identifiers. Your AI client connects directly to PubMed Central through this MCP, giving you immediate access to millions of open-access articles. You can ask your agent to find all studies that mention a specific gene, then pull the full text for three of them in structured JSON format. The system handles complex data mapping, automatically converting between identifiers like PMCID, PMID, and DOI so everything lines up perfectly. Whether you're tracking how much one paper influenced subsequent research by finding articles that cite it, or just needing a summary of key metadata from an open-access record, this MCP makes the entire process conversational. Because we host over 4,000 tools in our catalog, Vinkius ensures your agent can access PMC alongside everything else you need for serious academic work.

Core Capabilities

01 — Find and filter literature

Search the entire archive using advanced criteria like author name, publication date, or specific keywords.

02 — Extract full-text content

Retrieve the complete article text in structured JSON or XML format for immediate data analysis.

03 — Map scientific identifiers

Convert between different types of academic IDs (PMCID, PMID, DOI) to keep your research sources consistent.

04 — Analyze citation impact

Track the influence of a specific paper by finding all subsequent articles that cite it.

05 — Gather article summaries and details

Pull metadata, license information, or simple abstracts for multiple open-access records quickly.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to the MCP within Vinkius and provide your NCBI Tool Name and email address.

02 Your agent authenticates with the service, making sure all credentials are linked for high-volume querying.

03 You ask your AI client a question—like 'Find me articles citing X'—and the system returns structured data directly into your workspace.

The bottom line is that you get organized, deep scientific knowledge delivered through plain conversation, bypassing manual database work entirely.

Built For

This MCP targets researchers and clinical data specialists. It's for the academic who can't afford to spend hours manually sifting through PubMed's interface just to get a structured list of citations, or the scientist who needs to process dozens of open-access articles into a single dataset.

Academic Researcher

Uses the tool to quickly find relevant literature for their thesis, pulling full text using get_bioc_article and tracking impact with get_citing_articles.

Bioinformatics Scientist

Automates the collection of biomedical datasets by running searches via search_articles and harvesting metadata records using oai_pmh_request for large-scale analysis.

Healthcare Professional

Verifies clinical facts or reviews the latest studies by asking the agent to summarize key findings from an article summary, rather than reading through multiple abstracts.

What Changes When You Connect

01 Stop copying and pasting identifiers. Use convert_ids to instantly map PMCID, PMID, and DOI across your entire project dataset, keeping everything clean for analysis.

02 Ditch manual PDF downloading. With get_bioc_article, you pull the full text of an article directly into a structured JSON format that your agent can read, analyze, and summarize immediately.

03 Track academic influence effortlessly. Instead of guessing impact, use get_citing_articles to automatically find all the papers that built on a specific study, giving you a clear picture of its reach.

04 Gather massive metadata sets in bulk. Use oai_pmh_request or get_oa_record to harvest licensing info and file locations for dozens of open-access records simultaneously.

05 Cut down research time instantly. Instead of reading abstracts one by one, use search_articles to filter millions of papers based on precise criteria like 'gene X' published after 2020.

Real-World Applications

Mapping a literature gap

A researcher needs to prove that no one has studied the interaction between Protein A and Gene B in mammals. They use search_articles to find all relevant papers, then they run get_citing_articles on key publications to ensure every potential connection has been documented.

Processing a cohort dataset

A data scientist needs the full text and license information for 50 specific open-access articles. They use oai_pmh_request or get_oa_record to harvest all required metadata before using get_bioc_article on each one.

Fact-checking clinical claims

A medical professional needs instant verification of a drug's efficacy. They ask their agent to find articles about the drug and then use get_article_summary to pull quick, reliable metadata on the most recent clinical trials.

Building a structured knowledge graph

A bioinformatics team needs to build a database of relationships between genes. They run search_articles for gene pairs and then use get_bioc_article to extract the specific data points from the full text, making it machine-readable.

Patterns to Avoid

Using generic web searches

✗ AVOID

Searching Google for 'CRISPR gene editing' and hoping to find the correct full articles. This results in paywalls, news reports, and random links.

✓ INSTEAD

Use search_articles directly with this MCP. Specify authors or publication dates right away so your agent pulls only verifiable data from PubMed Central.

Ignoring ID consistency

✗ AVOID

Finding an article using a DOI, but then trying to use that identifier in another tool without converting it first.

✓ INSTEAD

Always run convert_ids when you are moving between different types of identifiers (like turning a DOI into a PMCID) to guarantee the data points match up.

Over-relying on abstracts

✗ AVOID

Assuming the abstract provides enough detail for quantitative analysis. You'll end up with qualitative, non-structured text.

✓ INSTEAD

For deep work, you must use get_bioc_article to pull the full JSON content. This structured data is what your agent needs to actually analyze.

The Right Fit

Use this MCP if your workflow requires access to massive, structured, and verifiable biomedical literature—the kind of data that lives behind academic identifiers. You need tools like get_bioc_article or get_citing_articles because you're performing synthesis, not just reading. Don't use this if you simply want general background information on a topic; for that, stick to standard web search. However, if your goal is to build a structured dataset of gene interactions from multiple sources, this MCP is essential. If you only

need basic citation counts and nothing more, then the get_article_summary tool might be enough, but remember it's always best practice to validate everything using convert_ids first.

Sifting through academic literature feels like archaeology.

Today, getting comprehensive research requires a painful workflow. You search PubMed manually, bookmarking dozens of links. Then, you have to copy the DOI from one tab and paste it into another tool just to get the PMCID for comparison. Finally, if you want the full text, you click through multiple paywalls or spend an hour downloading PDFs just to extract three key data points.

With this MCP, your agent handles all of that mess. You ask a single question—for example, 'Which papers cite X and what are their findings?'—and the system executes searches_articles, converts IDs, retrieves full text using get_bioc_article, and summarizes everything for you in one go.

Get Structured Data from PubMed Central with article summary.

The manual steps that disappear are the tedious cross-referencing of identifiers and the headache of extracting key facts from unstructured text. You no longer have to manually track whether a paper is open access or what its license demands for reuse.

Now, you treat academic research like any other structured data source: you ask for it, you get it clean, and you use it.

PubMed Central: 7 Core Academic Tools

Use these seven specialized tools to search for literature, convert identifiers, pull full text, and analyze citation networks from PubMed Central.

#	TOOL	DESCRIPTION
01	`get_bioc_article`	Retrieves a complete scientific paper's content in structured JSON or XML format via the BioC API.
02	`get_citing_articles`	Identifies and lists PMC articles that reference a specific PubMed ID, mapping its academic impact.
03	`convert_ids`	Converts article identifiers between PMCID, PMID, and DOI formats to ensure data consistency across sources.
04	`get_oa_record`	Finds crucial information like citation details, licensing terms, and file locations for open-access articles.
05	`oai_pmh_request`	Harvests structured metadata from the PMC OAI-PMH Service for batch data collection.
06	`search_articles`	Searches PubMed Central using complex queries involving keywords, authors, and date ranges to locate articles.
07	`get_article_summary`	Pulls brief metadata summaries for PMC articles, giving you key details without needing the full text.

See It in Action

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

U Search PubMed Central for recent articles about 'CRISPR gene editing' published in 2023.

I found several articles. Notable ones include PMC1012345 ('Advances in CRISPR...') and PMC1023456. Would you like a summary of the most relevant one?

U Get the full-text content of article PMC7840891 in JSON format.

I've retrieved the BioC JSON for PMC7840891. It contains sections for Introduction, Methods, Results, and Discussion. Which part should I analyze first?

U Convert the DOI 10.1038/s41586-020-2012-7 to a PMCID.

The DOI 10.1038/s41586-020-2012-7 maps to PMCID: PMC7095063 (and PMID: 32015508).

Frequently Asked Questions

01 **How do I find all the papers related to a specific study using PubMed Central MCP?**

You run the get_citing_articles tool. This finds articles that reference a given PMID, letting you track the scientific impact and lineage of research.

02 **Does PubMed Central MCP only handle one type of identifier?**

No. The convert_ids tool lets you reliably switch between PMCID, PMID, and DOI formats, keeping your data consistent no matter where the source ID came from.

03 **Can I get the full text in JSON format using PubMed Central MCP?**

Yes. You use get_bioc_article to retrieve the complete article content structured as BioC XML or JSON, making it immediately usable for analysis.

04 **What is the best way to collect metadata from multiple open-access articles?**

You can combine search_articles with get_oa_record. First find the list of article IDs, then run get_oa_record on those IDs to gather citation data and license info in bulk.

05 **Is oai_pmh_request better than getting a summary?**

They serve different purposes. Use get_article_summary for a quick, simple abstract. Use oai_pmh_request if you need to programmatically harvest large batches of structured metadata.

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 `"pubmed-central": { "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

PubMed Central 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

DOCUMENT INFORMATION

Generated	June 2026
MCP Server	PubMed Central MCP
Server ID	`019e38dc-f84b-7280-ae76-17ec71654566`
Platform	Vinkius Cloud for AI Agents
Endpoint	`https://edge.vinkius.com/{token}/mcp`

LICENSE & USAGE