

MCP SERVER

NO CODE

CLOUD HOSTED

MIT DBLP MCP

Map research networks and publication history.

MIT DBLP connects your AI client to the DBLP Computer Science Bibliography, giving you direct access to millions of academic papers. Use it to build detailed author profiles, map complex co-author networks, and search across major CS conferences like NeurIPS, ICML, and SIGMOD.

A+ Quality Score 98.33/100

academic-research

bibliography

computer-science

citation-network

data-indexing

search-engine



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.

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.

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.

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

Encryption at rest

Ed25519

PKI vault signatures

24h TTL

Ephemeral session keys

V8 Isolate

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.

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.

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.

06 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot 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.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

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.

MIT DBLP MCP

16 tools available

Cloud-hosted on Vinkius

This MCP lets your agent dive deep into computer science research history using the DBLP Bibliography index. You can ask it to find all papers related to a specific topic at top venues or track how an author's work has evolved over decades. Instead of manually navigating conference websites, you tell your AI client what you need—for example, 'Show me all database system papers from SIGMOD in the last five years.' It gathers the metadata and citation patterns instantly. Whether you are writing a literature review or tracking a candidate's publication record, this MCP structures that vast amount of academic data into usable insights. Through Vinkius, your AI client gets access to this massive catalog, letting you focus on analyzing the findings instead of curating the data.

Core Capabilities

01 — Identify Author Profiles

Fetch a complete publication history and profile for any researcher using their DBLP ID.

02 — Track Collaboration Networks

Determine an author's co-authors, ranked by how often they have published together.

03 — Analyze Research Productivity

Get detailed publication statistics for any author, showing trends and total output counts.

04 — Search Major Conferences/Journals

Find all papers published at specific academic venues or conferences (e.g., NeurIPS 2024).

05 — Query Specific CS Domains

Filter searches to target niche areas, like AI/ML, systems, theory, or database papers.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/mit-dblp — connect your AI agent in three steps.

- 01 You tell your agent the scope of the research you need, specifying an author's name and a time period.
- 02 The MCP uses that context to search DBLP for specific records, gathering titles, co-authors, and citation data.
- 03 Your agent returns structured lists and statistics detailing the publication network and trends.

The bottom line is you get machine-readable academic metadata without having to manually query multiple databases or websites.

Built For

This MCP is essential for PhD students stuck in literature review purgatory, faculty tracking their departmental research output, and hiring committee members needing objective proof of a candidate's academic reach.

PhD Student

Using the tool to gather comprehensive related work by searching for specific papers or understanding the co-author network surrounding their topic.

Academic Faculty

Tracking departmental research output and publishing trends across multiple years to write grant proposals or annual reports.

Technical Recruiter / Hiring Committee

Evaluating a candidate's depth of work by checking their publication record, co-author connections, and contribution volume at top venues.

What Changes When You Connect

-
- 01 Track an author's full academic journey by calling `get_author_stats`, which provides metrics on total publications and year-over-year output trends. It gives you the objective data needed for tenure reviews or grant applications.

 - 02 Visualize collaboration patterns immediately. Use `get_coauthors` to rank a researcher's collaborators based on their joint publication count, quickly identifying key research groups in any field.

 - 03 Filter noise from millions of papers. Dedicated searches like `search_ai_papers` and `search_database_papers` let you narrow results down instantly to the specific domain you need (e.g., only ICML or only SIGMOD work).

 - 04 Conduct precise literature reviews by combining tools. You can use `get_venue_publications` combined with `search_in_venue` to find all relevant papers at a conference like NeurIPS, but only concerning 'large language models'.

 - 05 Gain deep context on any single piece of research using `get_publication`. This tool pulls every detail—DOI, abstract, key authors—for verification and immediate use in your report.
-

Real-World Applications

Reviewing a PhD candidate's background

A hiring committee member needs to assess if a candidate really worked on distributed systems. They ask their agent to run `search_systems_papers` for the last 5 years, then cross-reference those results with `get_coauthors` to see who the candidate repeatedly published with at top venues.

Tracking academic field evolution

A faculty member wants to write a review on AI progress. They ask their agent to run `search_ai_papers`, then use `search_by_year`, filtering results year by year. This allows them to document the precise shift in focus—from early deep learning concepts to modern transformer architectures.

Finding papers missed during research

A student is working on a specific topic at SIGMOD but can't find related work. They ask their agent to execute `search_database_papers`, then use `search_in_venue` with the specific conference abbreviation and keyword to pull up every relevant paper.

Mapping an author's reach

A researcher wants to know which venues an established colleague is publishing at. They ask their agent to first `get_author`, then use `get_coauthors` on the result, and finally run `search_publications` using the co-author's name to map out the entire network.

Patterns to Avoid

Treating it like a general web search

✗ AVOID

Asking your agent to 'Find papers on AI.' This is too broad and will return millions of results with no structure.

✓ INSTEAD

Don't just ask for 'AI.' Use the dedicated tools. Start by running `search_ai_papers`, or better yet, use `search_in_venue` by specifying a key conference like NeurIPS alongside your topic query.

Relying on general keyword searches

✗ AVOID

Searching only by title keywords without checking the publication date or venue. You risk finding outdated or irrelevant work.

✓ INSTEAD

Always constrain your search first. Use `search_by_year` to limit the timeframe, and use `get_venue_publications` to ensure the paper is from a reputable edition of a conference.

Ignoring co-author context

✗ AVOID

Listing papers by an author without knowing who they worked with. It makes the work look isolated.

✓ INSTEAD

Before writing about the author, run `get_coauthors` to identify their core research group and then use that information to contextualize their primary contributions.

The Right Fit

Use this MCP if your problem is structured academic citation analysis. You need to move beyond simple keyword searching; you require metadata like co-author links, formal publication statistics (from `get_author_stats`), and domain filtering (like `search_systems_papers`). This tool excels when you need

quantitative evidence of research influence or collaboration patterns.

Do NOT use it if your goal is general knowledge retrieval—for example, 'What are the top 5 CS topics?' For that, a simple LLM chat interface is enough. Also, if you just need to find a paper by its DOI alone, while `get_publication` can handle this, other specialized academic databases might offer more immediate access. This MCP's strength is in *network analysis* and *structured data retrieval*, not raw content reading.

The headache of tracking research influence today

Right now, if you need to assess an author's true impact or map out a field's history, you spend hours hopping between databases: checking Google Scholar for titles, then manually cross-referencing those papers on individual conference websites (NeurIPS, ICML), and finally trying to piece together co-author relationships in a spreadsheet. It's slow, messy, and prone to human error.

With this MCP, you skip the clicking entirely. You ask your agent to analyze the author's entire network, getting structured data that shows their full publication count, year-over-year trends, and collaborators—all pulled together in one go.

Discovering Publication Data with `get_author_publications`

Manually retrieving the most recent papers for an author means visiting their profile page and clicking through paginated results, often missing crucial metadata or having to copy/paste titles and years into a separate review document.

Instead, simply asking your agent to run `get_author_publications` provides up to 40 fully structured records instantly. You don't just get the titles; you get all the associated metadata ready for immediate analysis.

MIT DBLP with 16 Tools

These tools allow your agent to perform highly specific academic searches, from tracking co-authors to retrieving publication details for millions of CS records.

#	TOOL	DESCRIPTION
01	<code>get_author</code>	Retrieves an author's complete profile using their unique DBLP identifier.
02	<code>get_author_publications</code>	Lists up to 40 of the most recent papers associated with a specific author's name.
03	<code>get_author_stats</code>	Calculates an author's overall research productivity and impact metrics.
04	<code>get_coauthors</code>	Returns a ranked list of collaborators, showing who worked with the researcher most often.
05	<code>get_publication</code>	Gets all metadata for a single paper using its unique DBLP key.
06	<code>get_venue</code>	Provides details on academic venues, including full names and types of conferences or journals.
07	<code>get_venue_publications</code>	Lists all papers published at a specific edition of a conference (e.g., NeurIPS 2024).
08	<code>search_ai_papers</code>	Searches for the latest research specifically in artificial intelligence and machine learning from top venues.
09	<code>search_authors</code>	Finds computer science authors, providing disambiguated names and profile links.
10	<code>search_by_year</code>	Filters publication searches to only show papers from a specific year.
11	<code>search_database_papers</code>	Searches for the latest research focused on database systems at major conferences and journals.
12	<code>search_in_venue</code>	Finds papers within a specific conference or journal by combining the venue name with a topic query.
13	<code>search_publications</code>	Searches across all major CS venues for titles, authors, and details from millions of publications.
14	<code>search_systems_papers</code>	Finds the latest research focused on computer systems at top academic conferences.
15	<code>search_theory_papers</code>	Searches for theoretical computer science papers from specialized venues.

#	TOOL	DESCRIPTION
16	search_venues	Returns a comprehensive list of academic conferences and journals available in the index.

See It in Action

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

U Find recent AI papers on large language models at NeurIPS



I've found recent NeurIPS papers on large language models, including work on scaling, alignment, and efficiency.

U Search for publications by Yoshua Bengio



I've retrieved the complete DBLP profile for Yoshua Bengio, showing over 600 publications across NeurIPS, ICML, JMLR, and other venues.

U Find the latest database systems papers from SIGMOD and VLDB



I've searched DBLP for recent SIGMOD and VLDB papers on database systems, including work on query optimization, distributed databases, and transaction processing.

Frequently Asked Questions

01 How do I use MIT DBLP MCP to find papers from a specific conference?

You should use `search_in_venue`. You combine the exact venue name (like 'ICML') with your topic query, and it returns only relevant papers for that event.

02 Can I check if an author is active in a field using `get_author_stats`?

Yes. The `get_author_stats` tool gives you key metrics like total publication counts and venue distribution, which helps confirm research activity over time.

03 What's the difference between search_publications and search_in_venue?

search_publications is a broad net, covering all major venues. Use search_in_venue when you want to narrow results down to one specific conference or journal edition.

04 How do I find my collaborators using MIT DBLP MCP?

Run the get_coauthors tool. It gives a ranked list of co-authors, which is essential for understanding who influenced the researcher most.

05 Do I need to know the DBLP key to use get_publication?







Yes, get_publication requires the unique DBLP key. This key can usually be found within a publication's URL or metadata record.

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 <code>"mit-dblp": { "url": "..." }</code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → 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

MIT DBLP 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](https://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 MIT DBLP. 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	MIT DBLP MCP
Server ID	019dea5f-ccd8-72ab-ad18-63b2ae9daa66
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/mit-dblp.