MCP SERVER | NO CODE | CLOUD HOSTED

CourtListener MCP for AI Agents

Analyze case law, dockets, and judicial records for legal research.

CourtListener delivers a specialized Model Context Protocol for AI agents that handles complex legal research. Your agent can search court opinions, audit dockets, and pull detailed judge information by simply asking questions in natural language. It turns hours of manual database searching into instant, conversational answers.

A+ Quality Score 100/100

legal-research | court-opinions | docket-tracking | case-law | judicial-data

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

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.

CourtListener MCP

10 tools available
Cloud-hosted on Vinkius

Managing legal research used to mean jumping between dozens of specialized portals, cross-referencing citation data, and hoping you didn't miss a key docket entry. Now, your AI agent handles the whole workflow. By connecting CourtListener to your agent—whether that's through Claude, Cursor, or any MCP-compatible client—you don't interact with technical forms; you talk to an expert assistant. You can ask your agent to find all recent opinions on a specific legal topic, retrieve detailed court dockets for ongoing cases, and even verify the financial disclosures of a judge involved in litigation. It makes complex case law research feel like having a full-time paralegal sitting next to you. This entire capability is managed through Vinkius, making it easy to connect your AI workflow without needing deep legal tech expertise.

Core Capabilities

01 — Search and retrieve court opinions

Find specific legal opinions using keywords, retrieving metadata like the court name and date filed.

02 — Monitor ongoing litigation via dockets

Query court dockets to stay updated on active case filings and procedural movements in real time.

03 — Audit judicial records and finances

Get detailed profiles for judges, including information about their service dates and financial disclosures.

04 — Map legal precedents via citations

List the specific citations associated with an opinion so you can trace its place within the legal network.

05 — Identify relevant courts and jurisdictions

Retrieve a list of all available courts, ensuring your research is accurate for the correct jurisdiction.

One Click on Vinkius — From Prompt to Execution

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

01 First, subscribe to this MCP on Vinkius.

02 Second, input your CourtListener API key into your AI client's settings.

03 Third, ask your agent a natural language question—like 'What are the recent opinions regarding copyright fair use?'—and it handles the rest.

The bottom line is, you tell your agent what legal information you need, and it pulls structured data from court databases directly into your chat window.

Built For

Legal professionals, investigative journalists, and law students who struggle with the sheer volume of raw legal text. If manual case history research takes up hours of your day, this MCP is built for you.

Corporate Counsel

You need to track how specific regulations are being interpreted in recent court rulings. You use the agent to search for opinions and audit dockets related to compliance issues.

Investigative Journalist

Your job is verifying facts on powerful figures, so you run background checks using the agent to retrieve judge profiles and financial disclosures.

Academic Researcher

You need to build a comprehensive understanding of case law. You use the agent to list citations for key opinions, mapping out legal precedents automatically.

What Changes When You Connect

01 Instantly audit court dockets: Use `search_dockets` to monitor ongoing litigation and see exactly what cases are actively filed in a given jurisdiction without manual site visits.

02 Deep dive into judge history: Pull profiles using `list_judges` or `get_judge`, and then verify their financial records with `list_financial_disclosures`. This keeps your research objective and grounded.

03 Understand legal context fully: Don't just read an opinion. Use `list_citations` to map out the entire network of precedents that influenced the ruling, giving you full background knowledge.

04 Efficient case gathering: Start broad by using `search_opinions`, then narrow your focus instantly with `get_opinion` to retrieve all necessary details for a single case file.

05 Jurisdictional certainty: Before starting any search, use `list_courts` or `get_court` to confirm the exact court and rules governing your research area. You never want to mix up jurisdictions.

Real-World Applications

Tracking a specific case's progression

A paralegal needs to know if 'Acme Corp' filed anything new this month. They ask their agent, and it uses `search_dockets` to find the most recent filings, returning case numbers and filing dates immediately.

Verifying legal standing for a client

A corporate counsel needs to know if a specific law has been challenged recently. They ask their agent to use `search_opinions` with keywords like 'statute X' and receive the top 5 most relevant, recent opinions.

Building a background file on an opponent

A journalist needs deep context on a judge. They ask their agent to use `get_judge` followed by `list_financial_disclosures`, building a comprehensive profile that verifies the judge's professional history.

Patterns to Avoid

Confusing opinions with dockets

✗ AVOID

A user searches for a case name and only gets general court information. They assume they've found everything, but miss the active filings.

✓ INSTEAD

To find what's happening *right now*, use `search_dockets`. If you need to read the actual ruling text, make sure you run `search_opinions` first and then retrieve details using `get_opinion`.

Missing jurisdiction context

✗ AVOID

A researcher finds an opinion but can't confirm if it applies to their state. They spend hours cross-referencing manual court lists.

✓ INSTEAD

Always start by using `list_courts` or `get_court`. This verifies the rules and metadata for your specific jurisdiction before you run any search.

Overlooking related legal context

✗ AVOID

A user finds a key opinion but doesn't know which other cases relied on it. The research feels incomplete.

✓ INSTEAD

After finding an opinion, immediately use `list_citations`. This tool shows you the full network of legal precedents that influenced or referenced that specific ruling.

The Right Fit

Use this MCP if your work depends on structured data from court systems. You need to find *what* was filed (dockets), *who* made the decision (judges, financials), and *why* it matters (opinions and citations). Don't use this if you are just looking for general legal summaries or articles; those require different types of sources. If your goal is simply to list all available courts without running a search, `list_courts` handles that efficiently. But if you need deep-dive comparison across multiple domains—like matching a judge's financial status against their rulings—this tool gives you the specific mechanisms like `get_judge` and `search_opinions` to link them together.

CourtListener MCP: Solving Manual Case Law Research with Opinions

Today, conducting deep legal research means juggling multiple websites. You find an opinion, but you have to manually copy the citation and then navigate to a separate docket portal just to see if that case has been updated since the ruling was filed. It's tedious, error-prone work that eats up days.

With this MCP, your agent handles it all. Ask for opinions on a topic, and you get not only the text but also immediate access to its citations and the ability to cross-reference judge details. You stop copying data and start receiving actionable intelligence.

CourtListener MCP: Auditing Judge Backgrounds and Dockets

Before making any legal assumption, you have to verify who made the ruling. This means manually checking judge profiles, hunting down their service dates, and trying to locate public financial disclosures—a process that takes weeks.

Now, your agent runs background checks in seconds. You can use `list_judges` to get a roster, then deep-dive into any profile using `get_judge`, followed by `list_financial_disclosures`. The data is instantly synthesized for you.

CourtListener: 10 Tools for Legal Research & Case Law Analysis

These tools let your agent perform specific actions like searching dockets, getting judge profiles, or listing citations across legal databases.

#	TOOL	DESCRIPTION
01	`get_judge`	Pulls specific biographical and professional details about a named judge.
02	`list_citations`	Provides all the legal citations tied to a particular court opinion.
03	`list_financial_disclosures`	Retrieves public records of a judge's financial filings and assets.
04	`list_judges`	Generates an alphabetized list of all judges currently recorded in the system.
05	`list_opinions`	Lists available legal opinions, providing basic details for further investigation.
06	`get_opinion`	Fetches the full text and detailed metadata for a specific court opinion using its identifier.
07	`search_dockets`	Searches current court dockets to find active case filings based on names or dates.
08	`search_opinions`	Searches the full database of legal opinions using keywords, date ranges, and circuit identifiers.
09	`get_court`	Retrieves specific details about a court system or jurisdiction by name.
10	`list_courts`	Lists all known courts and their respective metadata to confirm the correct legal jurisdiction.

See It in Action

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

U I need to know all the recent cases involving intellectual property and where they were filed.

🔍 **Litigation Report: Intellectual Property (Last 90 Days)**

- **Circuit:** 2nd Circuit
- **Case Name:** InnovateTech v. GlobalMedia
- **Status:** Active Docket Filing (Filed: June 1, 2024)
- **Summary:** Dispute over digital media transformation rights.

*

Top Opinions Found (Keyword Match):

Opinion	Court	Date Filed	Citation
Smith v. Jones	9th Circuit	May 15, 2024	[Citation details]
Copyright Corp	2nd Circuit	Apr 30, 2024	[Citation details]

U Can you pull up the profile and financial records for Judge Smith?

👩‍⚖️ **Judge Profile: Evelyn Smith**

- **Jurisdiction:** Federal Circuit Court
- **Service Start:** 2015
- **Specialty:** Patent Law, IP Dispute Resolution.

💰 **Financial Disclosures (Last Filing):**

- **Date:** Q4 2023
- **Assets Reported:** $1.2M - $3.5M
- **Discrepancies Noted:** None detected by agent review.

(This data is sourced from the official court records.)

U Show me opinions about 'data privacy' and list all potential precedents.

📄 **Data Privacy Opinions Report**

Found 4 key opinions matching 'data privacy' in the last year.

Key Result: The most frequently cited opinion is *Doe v. State* (2023).

📚 **Related Precedents for Doe v. State:**

1. *Smith et al.* (Defining personal data scope)
2. *Federal Comm. Board v. Telecom* (Establishing industry standards)
3. *Garcia v. Agency X* (Early ruling on consent forms)

Frequently Asked Questions

01 **How does CourtListener help me search for case law and opinions?**

CourtListener lets your AI agent perform complex searches across massive databases of court opinions using natural language. Instead of writing complicated search queries, you just ask what you need—for example, 'What are the recent rulings on data privacy?'.

02 **Can I use CourtListener to audit judge credentials?**

Yes. You can pull detailed profiles for specific judges and check their professional history and financial disclosures. This is critical for verifying who has been making legal decisions in your area of research.

03 **What if I need to track a case's current status?**

You can use CourtListener to search active court dockets. Your agent will pull the most recent filings and procedural updates for a given case, letting you monitor litigation without visiting multiple websites.

04 **Is CourtListener useful for tracking legal precedents?**

Absolutely. If your agent finds an important opinion, it can instantly list all the citations associated with that ruling. This shows you exactly which previous cases set the groundwork for the current law.

05 What kind of courts does CourtListener cover?

The MCP allows you to list and get details for various court systems, helping you confirm your jurisdiction before you start any search. This ensures all your research is legally accurate from the outset.

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 `"courtlistener": { "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

CourtListener 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 CourtListener. 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	CourtListener MCP
Server ID	`019d842a-9548-719d-ae13-505e890a3777`
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/courtlistener.