MCP SERVER | NO CODE | CLOUD HOSTED

Stanford bioRxiv MCP

Track Science From Preprint to Publication

Stanford bioRxiv connects your AI agent directly to the world's leading open-access preprints for biology (bioRxiv) and health sciences (medRxiv). Use this MCP to search, track revisions, and find cutting-edge research papers months before they appear in peer-reviewed journals. It lets you explore deep scientific categories like genomics, oncology, and neuroscience by date range or author institution.

A+ Quality Score 98.33/100

biorxiv | medrxiv | preprints | biology | health-sciences | neuroscience

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

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.

Stanford bioRxiv MCP

16 tools available
Cloud-hosted on Vinkius

Need the absolute latest science? This MCP connects your AI agent straight into bioRxiv and medRxiv—the major repositories where researchers post preprints for biology and health sciences. Instead of waiting months for journal publication, you get access to findings the day they are shared.

It’s built for anyone who needs to know what's happening at the cutting edge of research. You can filter papers by specific disciplines—think cancer immunology or developmental biology—or track how a single manuscript changes over time by checking its full revision history. Furthermore, you don't have to guess if a preprint made it into print; the MCP tracks which preprints eventually become published in peer-reviewed journals. By connecting this functionality through Vinkius, your AI client can act as an instant research librarian, pulling together complex metadata like DOI numbers and author affiliations across massive datasets of scientific findings.

Core Capabilities

01 — Retrieve full paper details by DOI

Find a preprint's title, authors, abstract, and institutional affiliation using its unique Digital Object Identifier.

02 — Monitor manuscript revision history

Trace the evolution of an idea by retrieving every version a preprint has gone through over time.

03 — Track publication status

Determine if a specific preprint has been accepted and published in a formal, peer-reviewed journal.

04 — Browse the latest submissions

Get immediate access to the most recent papers submitted across both bioRxiv and medRxiv categories.

05 — Filter by scientific discipline

Narrow down results instantly using specialized subject feeds, such as genomics, epidemiology, or cell biology.

One Click on Vinkius — From Prompt to Execution

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

01 Tell your agent what you need to research—for example, 'Find the latest preprints on cancer immunology from last month.'

02 The MCP executes targeted searches across bioRxiv and medRxiv using the specified filters (date range, category, etc.).

03 You get back a structured list of results containing abstracts, author details, version history, and publication tracking metadata.

The bottom line is you skip the manual searching through multiple websites and just ask your agent to pull the data for you.

Built For

Anyone who works with rapidly changing scientific fields. This MCP is crucial for biologists, medical researchers tracking emerging outbreaks, PhD students needing fresh literature reviews, or science journalists covering breaking discoveries.

Biologist

Uses the MCP to stay ahead of journal cycles by finding early-stage preprints in niche areas like developmental biology before competitors see them.

Medical Researcher

Monitors medRxiv for clinical and public health outbreaks, using tools to track which initial findings eventually get peer-reviewed status.

Science Journalist

Quickly pulls the latest submissions from bioRxiv or uses specific category filters (like epidemiology) to report on breaking scientific trends.

What Changes When You Connect

01 You stop waiting for journal publishing cycles. With the `get_recent_biorxiv` and `get_recent_medrxiv` tools, you get immediate access to research findings the day they're shared.

02 You don't have to piece together fragmented data points. Use `search_by_category` or `search_genomics` to instantly filter thousands of papers by a precise discipline like 'cell biology'.

03 You always know if that groundbreaking paper you read is final. The MCP's publication tracking tools let you check both the initial preprint status and use `get_published_version` to find the peer-reviewed source.

04 Tracking research evolution becomes simple. Instead of guessing how a concept changed, `get_preprint_versions` shows every revision history for maximum data integrity.

05 You can narrow massive datasets down by origin. Use `search_by_institution` to see what specific organizations are publishing the most cutting-edge work in a given timeframe.

Real-World Applications

Monitoring an emerging viral outbreak

A public health researcher uses the MCP with `search_epidemiology` and `get_recent_medrxiv`. They monitor for new papers on transmission vectors, allowing them to advise policy makers before any journal publication occurs.

Reviewing novel cancer targets

A PhD student needs the latest data. Using `search_cancer` and filtering by date range via `search_biorxiv`, they pull 50 relevant preprints, saving weeks of manual database trawling.

Verifying scientific claims

A journalist reads a promising preprint. They immediately run `get_published_tracking` on the DOI to verify if it has been submitted or accepted by a peer-reviewed journal, adding credibility to their story.

Comparing research approaches

A lab director wants to see who's working in his field. They use `search_by_institution` to map out which universities are currently publishing the most active preprints in molecular biology.

Patterns to Avoid

Assuming all preprints are final

✗ AVOID

Reading an abstract and assuming its findings are peer-reviewed fact, citing it as absolute truth.

✓ INSTEAD

Always check the `get_published_tracking` status. If it hasn't been published, treat the data as preliminary research only.

Searching too broadly by keyword

✗ AVOID

Using a general search function without defining time limits or specific disciplines.

✓ INSTEAD

Use targeted tools like `search_by_category` (e.g., 'immunology') combined with `search_biorxiv` using precise date ranges.

Ignoring revision history

✗ AVOID

Using the details from a preprint that was revised six months ago, missing critical updates.

✓ INSTEAD

Always run `get_preprint_versions` to understand how the manuscript has evolved and if the latest version contains crucial corrections.

The Right Fit

Use this MCP when your priority is speed and cutting-edge information. If you need to see what's happening *right now* in biology or medicine, this tool provides that direct feed from researchers before the formal academic gauntlet. Don't use it if you absolutely must cite a paper as final, peer-reviewed fact; for that, rely on databases of published journals, not preprints. When verifying publication status, always start with `get_published_tracking`. If your goal is purely historical research—say, tracking how vaccine development ideas have changed over decades—then use the version tracking tools. But if you just need a general overview without specific filters, don't waste time; stick to dedicated category searches like 'genomics'.

The Manual Struggle of Finding Cutting-Edge Science

Today, tracking novel research is an exhausting process. You have to jump between multiple preprint sites, manually copy DOI numbers, and cross-reference dates just to figure out if a finding is new or already published. This means wasting hours sifting through thousands of abstracts that might be outdated or incomplete.

With this MCP, you don't do the heavy lifting anymore. You tell your agent what topic you need—say, neuroscience preprints from Q2—and it pulls everything together instantly. You get a structured list containing all the necessary metadata and links to the latest submissions.

Get Scientific Certainty with `get_published_tracking`

The biggest frustration is ambiguity: Is this finding preliminary, or has it passed rigorous review? You waste time checking multiple databases to verify the status of a DOI.

Now, running `get_published_tracking` tells you immediately if that preprint made it through the peer-review process. It gives you definitive answers about whether the work is published in a journal and when—no more guesswork.

Stanford bioRxiv: 15 Tools for Scientific Discovery

These tools let you systematically query vast scientific repositories, enabling precise searches across specific disciplines, date ranges, and author institutions.

#	TOOL	DESCRIPTION
01	`get_preprint`	Searches both bioRxiv and medRxiv. Returns title, authors, corresponding author and institution, date, version, category, abstract, and license. DOI format: "10.1101/2024.01.15.575123". Get preprint details by DOI
02	`get_preprint_versions`	Preprints on bioRxiv/medRxiv can be updated multiple times. This lets you see the full revision history and understand how a manuscript has evolved. Get all versions of a preprint to track revisions
03	`get_published_tracking`	Shows the preprint DOI, published DOI, journal name, and publication date. Essential for understanding the preprint-to-publication pipeline. Track which preprints have been published in journals
04	`get_published_version`	Returns the published DOI, journal citation, and publication date. Essential for finding the final, peer-reviewed version of a preprint you have read. Find the journal-published version of a preprint
05	`get_recent_biorxiv`	Default is 7 days. Essential for staying at the cutting edge of biological research — preprints appear here 6-12 months before peer-reviewed publication. Get the latest bioRxiv preprints
06	`get_recent_medrxiv`	Covers clinical medicine, epidemiology, public health, and health systems research. Critical for monitoring emerging health research before journal publication. Get the latest medRxiv preprints
07	`search_biorxiv`	The bioRxiv API returns preprints in batches of 100. Use the date interval format "YYYY-MM-DD/YYYY-MM-DD" (e.g. "2024-01-01/2024-01-31"). Use cursor for pagination (0, 100, 200, etc.). Browse bioRxiv preprints by date range
08	`search_by_category`	bioRxiv categories include: neuroscience, genomics, bioinformatics, cell_biology, cancer_biology, immunology, microbiology, molecular_biology, biochemistry, genetics, developmental_biology, evolutionary_biology, ecology, plant_biology, physiology, pharmacology, systems_biology, biophysics, synthetic_biology. medRxiv categories: epidemiology, infectious_diseases, public_and_global_health, health_systems, cardiovascular_medicine, oncology, psychiatry, neurology. Filter preprints by subject category
09	`search_by_institution`	Use this to explore what institutions are producing preprints in a given time period. Each preprint includes the corresponding author and their institutional affiliation. Browse preprints with author institution metadata

#	TOOL	DESCRIPTION
10	search_cancer	Covers tumor biology, oncogenomics, cancer immunology, drug resistance, and experimental therapeutics. Browse cancer biology preprints
11	search_cell_biology	Covers cell signaling, organelle biology, cytoskeleton, cell division, stem cells, and cellular mechanisms of disease. Browse cell biology preprints
12	search_epidemiology	Covers disease surveillance, outbreak analysis, population health, health policy, and clinical epidemiology. Critical for public health monitoring. Browse epidemiology and public health preprints
13	search_genomics	Covers genome sequencing, gene regulation, epigenomics, metagenomics, and computational genomics — core disciplines in modern biology. Browse genomics and bioinformatics preprints
14	search_immunology	Covers immune system research, host-pathogen interactions, vaccine development, autoimmune diseases, and immunotherapy. Browse immunology and microbiology preprints
15	search_medrxiv	medRxiv covers clinical research, epidemiology, public health, and health policy. Use interval "YYYY-MM-DD/YYYY-MM-DD" format. Results paginated in batches of 100. Browse medRxiv preprints by date range
16	search_neuroscience	Neuroscience is one of the largest and most active categories, covering brain research, neural circuits, cognitive science, and neurological disorders. Browse neuroscience preprints

See It in Action

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

U Show me the latest neuroscience preprints

I've retrieved the latest neuroscience preprints from bioRxiv. Recent submissions cover topics including neural circuits, brain imaging, cognitive neuroscience, and neurological disorders.

U Has preprint 10.1101/2024.01.15.575123 been published in a journal?

I've checked the publication tracking for this preprint. If published, I'll show the journal DOI, citation, and publication date.

U Find the latest genomics preprints from this week

I've retrieved this week's genomics preprints from bioRxiv, including recent work on single-cell sequencing, CRISPR screens, and population genetics.

Frequently Asked Questions

01 **How do I find the most recent preprints using Stanford bioRxiv MCP?**

You can use `get_recent_biorxiv` for general biological updates, or if you're focused on health, run `get_recent_medrxiv`. These tools pull the absolute latest submissions across their respective platforms.

02 **Can I check if a preprint was published using Stanford bioRxiv MCP?**

Yes, use `get_published_tracking` with the DOI. This tool tells you if the work has been accepted into a peer-reviewed journal and provides citation details.

03 **What is the best way to search for cancer research preprints?**

Use the dedicated `search_cancer` tool. This focuses your search specifically on tumor biology, oncogenomics, and related areas within both bioRxiv and medRxiv.

04 **How do I see if a preprint was updated?**

You must run `get_preprint_versions` on the DOI. This function provides the full revision history, letting you track how authors changed their data or conclusions over time.

05 **Which tool should I use to find papers from a specific date range?**

For biology, run `search_biorxiv` using the 'YYYY-MM-DD/YYYY-MM-DD' interval. For health sciences, use `search_medrxiv` with the same date format.

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 `"stanford-biorxiv": { "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

Stanford bioRxiv 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 Stanford bioRxiv. 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	Stanford bioRxiv MCP
Server ID	`019dea60-1c4b-7316-8020-a54bd9e04747`
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/stanford-biorxiv.