

MCP SERVER

NO CODE

CLOUD HOSTED

AI21 Studio MCP for AI Agents

Refining Text Structure and Grammar for Content Creation

AI21 Studio offers advanced text processing for AI agents, letting you summarize massive documents, polish tone with paraphrasing, or correct complex grammar errors using Jamba models. It's built for developers and content creators who need industrial-strength linguistic tools to work directly through their favorite AI client.

A+ Quality Score 100/100

llm-integration

text-summarization

natural-language-processing

grammar-correction

foundation-models



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.

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](https://vinkius.com) — connect your AI agent in under 60 seconds.

AI21 Studio MCP

7 tools available

Cloud-hosted on Vinkius

Need your agent to handle messy text? This MCP gives your AI client specialized language tools that go way beyond standard chat prompts. You can feed it a huge wall of text, and it will instantly produce accurate summaries or break the content down into distinct sentences based on complex grammar rules. Need to change the tone—from super casual to highly formal? The paraphrasing tool handles all those style shifts for you. Developers love that they can convert any block of text into numerical embeddings right inside their workflow, making Retrieval Augmented Generation (RAG) much cleaner. If you're using a Vinkius-compatible client like Cursor or Claude, you just connect this MCP and your agent gets access to the entire suite. You never have to copy-paste messy raw data into a separate editor again.

Core Capabilities

01 — Generate Conversational Text

Your agent can generate chat responses using advanced Jamba models for fluid, natural conversations.

03 — Rewrite Text by Style

You can rewrite any passage into different styles, including general, casual, formal, long, or short versions.

05 — Break Text into Sentences

It splits continuous text blocks into individual sentences, respecting complex grammatical boundaries like initials or titles.

02 — Condense Long Documents

It takes large bodies of text and produces accurate, concise summaries instantly, saving you hours of reading time.

04 — Correct Grammar and Spelling

The dedicated engine finds and fixes grammatical mistakes and spelling errors in your draft copy.

06 — Create Numerical Embeddings

You pass in texts and get high-dimensional numerical vectors needed for building advanced search and knowledge systems (RAG).

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/ai21-studio — connect your AI agent in three steps.

- 01 Connect your preferred AI client to the Vinkius catalog and subscribe to this MCP.
- 02 Authorize access using your AI21 Studio API Key within your client settings.
- 03 Instruct your agent via natural conversation, telling it exactly what text operation you need (e.g., 'Summarize this article' or 'Fix the grammar in these paragraphs').

The bottom line is that once connected, your agent handles the entire workflow—it sends the request to the tool and brings back the corrected or processed result.

Built For

This MCP targets content teams and technical developers. It's for the editor who spends too long manually polishing copy, or the engineer building complex knowledge retrieval systems that need perfect text inputs.

Content Editor

Uses paraphrasing and grammar correction to polish rough drafts of marketing copy until they sound professional.

Technical Writer

Runs segmentation on complex manuals or documentation to ensure every section is grammatically isolated for publication.

Data Scientist / Developer

Generates embeddings from user-submitted documents and segments them into reliable chunks for a custom RAG application.

What Changes When You Connect

- 01 Stop struggling with manual copy-editing. Use the dedicated `grammar_corrections` tool to catch subject-verb agreement issues and spelling mistakes your eye usually misses.

-
- 02** When you need to reduce 50 pages of research into a usable executive summary, the `summarize` capability does it instantly without losing core meaning.
-
- 03** Need to change the voice? The `paraphrase` tool lets you rewrite text for different audiences—from highly formal white papers to casual blog posts.
-
- 04** Building advanced knowledge systems requires clean data. Use `embed_texts` to convert any corpus of documents into usable vectors for RAG workflows.
-
- 05** Handling massive, unstructured reports? The `segmentation` tool splits continuous paragraphs into grammatically distinct sentences, ensuring reliable processing later on.
-

Real-World Applications

Cleaning Up Raw Research Data

A researcher collects 80 pages of transcribed interviews. Instead of reading everything, they ask their agent to run `'summarize'` and get a bulleted overview highlighting key themes from the entire corpus.

Building a Custom Knowledge Base

The engineering team feeds thousands of technical manuals into their agent. Running `'embed_texts'` converts these documents into search vectors, allowing the AI client to answer highly specific questions accurately.

Polishing Client Pitches for Tone

A marketing director drafts a pitch that sounds too casual. They use `'paraphrase'`, specifying the 'formal' style, and get an immediate rewrite ready for the C-suite.

Standardizing Policy Documentation

A compliance officer has legacy documents with inconsistent language. They run `'grammar_corrections'` across all files, ensuring every policy adheres to modern grammatical standards.

Patterns to Avoid

Using generic LLM prompts for polish

X AVOID

Asking the agent: 'Make this text sound better.' The result is vague, often sounding overly polite or just repeating filler words.

✓ INSTEAD

Instead, use ``paraphrase`` and specify a precise style like 'formal' or 'casual'. Or if the issue is structural, run ``segmentation`` first to isolate complex sentences for targeted correction.

Ignoring text boundaries

X AVOID

Feeding one giant block of mixed data (questions and answers) into a single prompt. The LLM gets confused about which part needs fixing.

✓ INSTEAD

First, use ``segmentation`` to cleanly separate the questions from the answers. Then process each segment individually with ``grammar_corrections`` for maximum reliability.

Forgetting data preparation

X AVOID

Trying to feed raw text into a vector database without cleaning it first. The embeddings are noisy and retrieval fails.

✓ INSTEAD

Always run the source material through ``grammar_corrections`` and potentially use ``summarize`` or ``segmentation`` *before* generating embeddings with ``embed_texts``. Clean input equals clean output.

The Right Fit

Use this MCP if your core need is advanced linguistic manipulation: summarization, style shifting, or granular error correction. You must have a workflow where the structure or tone of text is critical—think technical documentation, legal briefs, or marketing copy. Don't use it if you simply need general creative brainstorming; standard chat completion will work fine for that. However, if your primary goal is *only* basic paraphrasing and you don't care about grammatical precision, a simpler, single-purpose text tool might suffice. If you are building a complex application relying on knowledge retrieval (RAG), this MCP is mandatory because the `embed_texts` capability provides the necessary vector preparation that generic tools lack.

AI21 Studio Text Processing for Content Editing

Content editors spend hours in Google Docs, manually checking style guides and fixing run-on sentences. They copy large chunks of draft copy into a new window, running through multiple grammar checkers just to polish the tone or find misplaced commas.

With this MCP, you simply tell your agent what's wrong—say, 'Rewrite this passage to sound like an industry expert.' The agent executes the `paraphrase` tool instantly and gives you polished text that hits the exact style mark you need. It's done in seconds.

AI21 Studio Text Processing for Knowledge Architecture

Developers often struggle with raw documentation. They copy massive, unstructured PDFs or websites into a database, which results in noisy data chunks that break when the AI tries to interpret them.

Using this MCP changes everything. You first use `segmentation` to reliably split text by grammatical boundaries. Then, you generate embeddings using `embed_texts`, creating clean, reliable knowledge pillars ready for any advanced search application.

AI21 Studio 7 Tools for Advanced Text Processing

Use these seven dedicated tools to analyze, correct, restructure, and embed complex natural language data through your AI client.

#	TOOL	DESCRIPTION
01	<code>chat_completion</code>	Generate conversational chat replies using AI21 Jamba models.
02	<code>text_completion</code>	Generate general text completions and continuations using AI21 Studio models.
03	<code>paraphrase</code>	Rewrite selected text passages into different styles like formal, casual, or short.
04	<code>summarize</code>	Create concise summaries from long blocks of text instantly.
05	<code>segmentation</code>	Split large chunks of text into separate sentences based on grammatical rules.
06	<code>grammar_corrections</code>	Perform detailed checks and corrections for grammar and spelling errors in provided text.
07	<code>embed_texts</code>	Convert arrays of texts into numerical embeddings, specifying if they are queries or documents.

See It in Action

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

U I have this article draft. Can you fix all the grammar errors and make it sound more professional?



Grammar Check Results:

- **Original:** The team is very excited about the results.
- **Correction:** The team was highly enthusiastic regarding the findings.

We found 3 issues, mostly relating to subject-verb agreement. Here's the fixed draft:

> Our dedicated staff are committed to ensuring that all future projects exceed expectations.

U Summarize this huge technical white paper for me. I just need five bullet points.



Summary: Key Takeaways

- The primary shift in the sector is toward decentralized ledger technology.
- Implementation requires robust API integration across multiple endpoints.
- Future scalability depends on optimizing data chunking and vector storage.
- Cost analysis suggests a 15% reduction in operational overhead within two years.
- Adoption rate projections show rapid growth over the next fiscal cycle.

U I need to break this long paragraph into individual, clean sentences.



Sentence Segmentation:

The market analysis shows strong growth potential. This trend was evident in Q3 reports. Furthermore, consumer adoption rates are accelerating globally. We recommend a pivot toward sustainable materials.

Frequently Asked Questions

01 How does the AI21 Studio MCP help me with grammar and spelling?

It uses a dedicated correction engine to find errors that general chat models often miss. You just point it at your text, and it returns a perfectly cleaned version, saving you manual proofreading time.

02 Can I use the AI21 Studio MCP to shorten large documents?

Yes, the summarization tool handles massive texts. It pulls out the core concepts and gives you an accurate summary without losing the most critical details for review.

03 Is AI21 Studio useful for developers building knowledge systems?

Absolutely. You use the `embed_texts` capability to turn your documents into search vectors, which is the foundation of any reliable, custom knowledge retrieval application.

04 What if I need to change the tone of my writing using AI21 Studio MCP?

The paraphrasing tool lets you control the voice precisely. You can switch a piece from 'casual' blog talk to 'highly formal' policy language with one command.

05 Is this better than just asking my agent for help? (AI21 Studio MCP)







Yes. This MCP connects you to specialized endpoints that perform single, high-precision tasks—like splitting sentences or correcting grammar—which is far more reliable and accurate than a general chat response.

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>"ai21-studio": { "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

AI21 Studio 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 AI21 Studio. 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	AI21 Studio MCP
Server ID	019d754a-310b-7310-9e3e-26722627e7ef
Platform	Vinkius Cloud for AI Agents
Endpoint	<code>https://edge.vinkius.com/{token}/mcp</code>

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/ai21-studio.