

MCP SERVER

NO CODE

CLOUD HOSTED

Jasper MCP

Generate campaigns with brand-consistent voice.

Jasper MCP gives your AI agent direct access to Jasper's content engine. Generate entire marketing campaigns—from blog posts and email sequences to ad copy and social captions—all while enforcing specific brand voices and styles. It lets you run templates, ground content in company knowledge, and orchestrate multi-channel creative output without ever opening the Jasper app.

A+ Quality Score 100/100

generative-ai

content-marketing

copywriting

brand-voice

campaign-orchestration



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 — connect your AI agent in under 60 seconds.

Jasper MCP

20 tools available

Cloud-hosted on Vinkius

This MCP connects your AI agent directly to Jasper's full suite of marketing tools. You can generate highly varied copy—like ad headlines, social posts, or email newsletters—and ensure it sounds exactly like your brand. The system manages specific style guides and configured voices so that every piece of content maintains a consistent tone, whether it's casual for Instagram or professional for an executive report. If you need to build complex campaigns, you can pass a single brief to generate multiple variations (e.g., ads, emails, and social posts) simultaneously. Because this MCP is hosted on Vinkius, your agent doesn't just get Jasper; it gains access to a catalog of thousands of specialized tools, letting you connect content generation alongside CRM data or analytics pipelines for one complete workflow.

Core Capabilities

01 — Generate multi-format copy

Create full content assets like blog posts, ad headlines, and social media captions from a single prompt.

02 — Enforce brand voice and style

Apply pre-configured brand voices or specific style guidelines to guarantee consistent tone across all generated material.

03 — Execute structured templates

Run specific, proven content formats—like the AIDA or PAS frameworks—by selecting and filling out defined templates.

04 — Ground copy in company knowledge

Ensure generated text is accurate by feeding it internal documents or specialized knowledge bases for reference.

One Click on Vinkius — From Prompt to Execution

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

- 01 First, subscribe to the Jasper integration on the Vinkius marketplace.
- 02 Next, generate an API token from your Jasper account settings and provide that key to your AI agent.
- 03 Finally, prompt your agent with a content request. It will use available tools—like running templates or applying brand voices—to return finished copy.

The bottom line is, you tell your AI client what marketing material you need, and it handles the complex process of styling, structuring, and generating that copy using Jasper's engine.

Built For

This MCP is built for content creators whose job involves constant context switching. It serves Copywriters who hate checking brand guidelines manually, Marketing Directors who need to scale campaigns quickly, and SEO Specialists who must constantly ground content in the latest research.

Content Marketer

Uses this MCP to take a single product announcement and generate optimized copy variations for email, LinkedIn, and Twitter simultaneously.

SEO Specialist

Runs the search knowledge tools to pull specific data points from internal documents, then uses that context to write new blog post drafts.

Campaign Manager

Develops full campaign sequences by running templates (like a 5-part email drip) and ensuring the tone matches the overall brand voice.

What Changes When You Connect

- 01 Stop context switching. Generate channel-specific copy—from ad headlines to social posts—and entire campaign variations in one request, eliminating the need to open multiple tools.

-
- 02 Maintain perfect brand tone. You can explicitly apply configured brand voices and style guides through tools like `list_voices` before running a command, ensuring every word aligns with your guidelines.

 - 03 Scale content creation efficiently. Use structured templates (`run_template`) for predictable outputs, such as email sequences or landing page copy, without manual drafting.

 - 04 Write smarter, not harder. When you need to reference company data, use tools like `search_knowledge` to pull facts from your internal knowledge base directly into the generated draft.

 - 05 Target specific groups. The MCP lets you define and apply content based on defined audiences using tools like `list_audiences` , making sure every piece hits the right person.
-

Real-World Applications

Launching a new SaaS feature

The Campaign Manager needs to launch a product update. They ask their agent for three deliverables: 1) A LinkedIn post, 2) Three Google Ad headlines, and 3) An email subject line. The agent uses the MCP's orchestration tools to generate all variations simultaneously, ensuring they all match the 'Excited/Enterprise' brand voice.

Developing A/B test copy

The Growth Team needs ad copy variations for a Black Friday sale. They prompt their agent to generate five headline options using `run_command`, asking it to simulate different tones (urgent, benefit-driven, discount-focused) for immediate testing.

Writing an SEO blog post

An SEO Specialist is writing about a new industry trend. Instead of researching manually, they ask their agent to use `search_knowledge` against internal research documents. The resulting content draft is automatically grounded in accurate company data, making the article credible and fact-checked.

Creating structured email flows

A Content Strategist needs a 3-part welcome sequence. They use the `list_templates` tool to find the 'Nurture Email' template and then use `run_task` to populate it with specific audience data, guaranteeing structural consistency across all emails.

Patterns to Avoid

Using Jasper as a simple prompt box

X AVOID

A user writes: 'Write me some copy about our product.' The output is generic because the AI has no context or style enforcement.

✓ INSTEAD

Instead, ask your agent to use ``run_command`` and explicitly provide both the brand voice ID and a knowledge item ID. This forces the tool to generate content that is both on-brand and factually accurate.

Ignoring defined style guides

X AVOID

The copy reads differently across different channels, sounding like it was written by three separate people.

✓ INSTEAD

Before generating any content, use ``list_styles`` to select the correct guide and pass its ID to your prompt. This ensures every piece of writing adheres to the same rules.

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

X AVOID

✓ INSTEAD

20 Tools Available

#	TOOL	DESCRIPTION
01	<code>get_usage</code>	Retrieves API usage statistics for defined date ranges and granularities.
02	<code>list_templates</code>	Lists all available content templates, showing their required inputs for execution.
03	<code>run_template</code>	Runs a specific content template by passing the necessary input values in JSON format.
04	<code>run_command</code>	Generates AI marketing content using a command prompt, allowing you to specify brand voices or context for grounding.
05	<code>list_voices</code>	Retrieves a list of stored brand voice profiles available for consistent content generation.
06	<code>get_voice</code>	Fetches specific details about an existing brand voice profile.
07	<code>list_knowledge</code>	Lists the knowledge items available to ground content generation in company facts.
08	<code>list_styles</code>	Provides a list of stored style guides that dictate tone and writing rules.
09	<code>list_documents</code>	Lists workspace documents, supporting filtering by name, status, or date range.
10	<code>get_document</code>	Retrieves detailed information about a specific document in the workspace.
11	<code>list_tasks</code>	Lists available agent tasks, showing the necessary context items required for execution.
12	<code>get_task</code>	Retrieves specific details and input requirements for an agent task.
13	<code>run_task</code>	Executes a defined agent task by passing necessary context items as a JSON array.
14	<code>list_users</code>	Retrieves a list of all users within the connected workspace.
15	<code>get_template</code>	Gets detailed information about a specific content template.
16	<code>get_knowledge_item</code>	Retrieves the full details of an individual knowledge item for context grounding.
17	<code>search_knowledge</code>	Searches the entire knowledge base by meaning and context, rather than just keywords.

#	TOOL	DESCRIPTION
18	<code>get_style</code>	Retrieves comprehensive details regarding a specific style guide for content generation.
19	<code>list_audiences</code>	Lists the predefined target audiences available to tailor generated content toward.
20	<code>get_audience</code>	Retrieves detailed information about a specific target audience profile.

See It in Action

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

U Write 3 Google Ads headlines for our SaaS product launch.



Here are 3 headline variations using your brand voice:

1. 'Scale Your Operations — Not Your Headcount'
2. 'Enterprise Automation, Startup Speed'
3. 'From Manual to Magical in 5 Minutes'

All headlines are under 30 characters and optimized for click-through.

U Generate an email subject line for our Black Friday campaign.



Here are 5 subject line options:

1. ' ⚡ 48 Hours Only: 40% Off Everything'
2. 'Your Biggest Savings of the Year Start Now'
3. 'Black Friday Alert: Deals Your Team Will Thank You For'
4. 'Don't Miss It — Our Best Price Ever'
5. 'Clock's Ticking: Black Friday Ends at Midnight'

U Create a LinkedIn post announcing our new AI feature.



Here's a LinkedIn post using your brand voice:

🚀 Big news from our engineering team.

We just shipped AI-powered workflow automation — and it's already changing how our customers operate.

What took 4 hours now takes 4 minutes.

No complex setup. No code. Just results.

Early adopters are seeing 10x efficiency gains in their first week.

Try it free → [link]

#AI #ProductLaunch #Automation

Frequently Asked Questions

01 How do I get my Jasper API token?

You need a **Jasper Business plan** to access the API. Log in at **app.jasper.ai**, go to **Settings → Dev Tools → API Tokens** (direct link: app.jasper.ai/settings/dev-tools/tokens). Click **Generate** to create a new token. Copy it immediately and paste it below. Only users with **Admin** or **Developer** roles can manage API tokens. If you don't see the Dev Tools option, verify your account is on the Business plan or contact your Jasper account admin.

02 Can my AI agent write ad copy in my brand's voice?

Yes. Jasper's Brand Voice feature ensures all generated content matches your company's tone, vocabulary, and style guidelines. Your AI agent accesses your configured brand voices and applies them automatically — so copy for Google Ads, LinkedIn posts, and email campaigns all sound authentically like your brand.

03 What if I need variations for A/B testing?

Just tell your AI agent 'Generate 5 headline variations for my Google Ads campaign.' Jasper produces multiple distinct alternatives per request, each maintaining your brand voice but varying the angle, hook, or call-to-action — giving your growth team ready-to-test creative in seconds.

04 Is the Jasper Business plan required?







Yes. API access is exclusive to the Jasper Business plan. If you're currently on Creator or Teams, you'll need to upgrade. Contact the Jasper sales team through jasper.ai to discuss Enterprise pricing — most marketing teams find the ROI justifies the investment within the first month of automated content production.

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>"jasper": { "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

Jasper 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 Jasper. 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	Jasper MCP
Server ID	019d75bc-82d8-715c-84d2-6f72a61fd46e
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/jasper.