

MCP SERVER

NO CODE

CLOUD HOSTED

ZenQuotes API MCP

Access daily, random, or batch quote content instantly.

ZenQuotes API delivers high-quality inspirational content directly to your AI client. Retrieve random quotes, get the official quote of the day, or pull batches of up to 50 quotes for deep research. It turns complex content sourcing into a simple conversation with your agent.

A+ Quality Score 100/100

inspirational-content

quote-api

motivational-data

content-curation

data-retrieval



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.

ZenQuotes API MCP

4 tools available

Cloud-hosted on Vinkius

Need reliable, fresh motivational data for an app or blog? This MCP connects your AI agent directly to ZenQuotes.io's library of inspirational quotes. Instead of manually visiting quote websites and copying text, you simply ask your agent what you need—a random boost, today's featured thought, or a thematic batch audit. Your agent handles the connection, fetching clean data instantly. You can run large-scale content research without touching any quote portal. Whether you're building a mindfulness app or just drafting blog posts, this service makes sure your content is always uplifting and perfectly formatted for immediate use in your workflow. Connecting it through Vinkius means that no matter which AI client you prefer, the ability to source motivational intelligence is right there, waiting for your command.

Core Capabilities

01 — Check API Health

You confirm if the quote service is running correctly before starting a workflow.

02 — Retrieve Random Quotes

Your agent fetches an immediate, random inspirational quote for quick content insertion.

03 — Get Daily Quote

You pull the official 'Quote of the Day' to maintain awareness of current trends in motivational content.

04 — Audit Large Batches

The system retrieves up to 50 quotes simultaneously, allowing you to audit a wide range of themes quickly.

One Click on Vinkius — From Prompt to Execution

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

- 01 Subscribe to this MCP on Vinkius. Since ZenQuotes is a free service, no API key setup is required.
- 02 Connect your preferred AI client (like Cursor or Claude) through the Vinkius platform.
- 03 Direct your agent to retrieve specific content, such as running ``get_zen_quote_of_the_day`` or calling ``get_random_zen_quote``. The data returns instantly for use.

The bottom line is that you get programmatic access to ZenQuotes' entire library of quotes without writing any code yourself.

Built For

This connector targets content creators, wellness coaches, and marketing operations leads. If your job involves constantly sourcing high-quality, fresh motivational text, this MCP eliminates the tedious manual copy/paste cycle.

Content Creator

You use ``get_zen_quotes_batch`` to pull dozens of quotes for a themed article or a social media campaign.

Mindfulness Coach

You run ``get_zen_quote_of_the_day`` at the start of your daily workshop prep, ensuring timely content delivery.

Marketing Operations Lead

You use random quote retrieval to quickly fill out placeholders or A/B test copy across multiple campaigns.

What Changes When You Connect

- 01 Stop manually searching for inspiration. Using `get_random_zen_quote` means your agent delivers a perfect quote in seconds, keeping your workflow moving fast.

-
- 02 Audit themes with volume: Instead of fetching quotes one by one, run `get_zen_quotes_batch` to get 50 quotes at once. This is essential for deep content research.

 - 03 Stay current and reliable: Use `get_zen_quote_of_the_day` every morning. You instantly know the industry's featured motivational thought without checking a website.

 - 04 Maintain data integrity by fetching unique metadata, allowing you to track authors and quote content reliably for attribution.

 - 05 Operational assurance is built-in: Run `check_api_status` whenever you start a major project to guarantee the source data is live.
-

Real-World Applications

Drafting a Blog Post Series

A content creator needs 20 quotes for an article on 'Wisdom.' Instead of opening a new tab and copying text repeatedly, they prompt their agent to use `get_zen_quotes_batch`. The agent pulls 50 quotes, and the creator filters them down to exactly what's needed in minutes.

Quick Social Media Content

A marketing team needs a rapid burst of quotes for Instagram stories. They ask their agent to call `get_random_zen_quote` three times in quick succession, getting fresh, varied content without any manual effort.

Daily App Feature Refresh

A mindfulness coach needs a fresh quote for their daily email newsletter. They ask their agent to run `get_zen_quote_of_the_day`. The agent instantly retrieves the official quote, which is then formatted and delivered straight into the draft email.

Researching Author Patterns

A researcher wants to see how many quotes are available from a specific period. They use `get_zen_quotes_batch` repeatedly and then process the metadata to build an audit of content distribution across authors.

Patterns to Avoid

Only checking for one quote.

X AVOID

You only ask your agent for a random quote when you actually need 15 quotes for a coming-of-age article. You end up making multiple, slow calls.

✓ INSTEAD

Don't rely on individual calls. Use ``get_zen_quotes_batch`` to pull the necessary volume of content in one go and then filter it using your AI client.

Ignoring service status checks.

X AVOID

You start a complex workflow right before a holiday or maintenance window, only to find that the quote source is down midway through processing.

✓ INSTEAD

Always run ``check_api_status`` first. This confirms the data pipeline is live and ready for your intensive content audit.

Forgetting structured retrieval.

X AVOID

You try to pull quotes from multiple different, unstructured sources because you didn't realize this MCP handles curated data. The results are messy.

✓ INSTEAD

Stick to the ZenQuotes source. Use ``get_zen_quotes_batch`` and rely on the standardized output format for guaranteed clean content.

The Right Fit

Use this if your primary need is sourcing high-volume, curated motivational text or quotes of various lengths. If you are building an app that needs daily inspiration, `get_zen_quote_of_the_day` is essential. Don't use this if you need complex data analysis like stock market predictions; for that, look at finance APIs. Similarly, if you need to analyze sentiment across thousands of user reviews, a dedicated NLP tool is better. However, if your goal is simply 'more content' or 'better copy,' the ZenQuotes API MCP gives you immediate access to reliable quote batches and single random picks.

The struggle of manual inspiration sourcing

Today, finding enough fresh, high-quality quotes for a blog or social campaign is a grind. You open Google, you click through three different quote websites, copy the text from one, check another site for attribution, and then paste it into your content draft. It's slow, manual, and always leaves you questioning if you missed a better source.

With this MCP, that entire process vanishes. Your AI agent handles the sourcing. You simply tell it what kind of quote you need—a random one or a batch of fifty for theme auditing—and the clean content appears instantly. The only thing left is writing your article.

ZenQuotes API MCP: Reliable Quote Data Access

No more opening multiple tabs or dealing with varying formatting across different quote sites. You never have to manually check the status of an external service again either.

This MCP delivers structured, reliable data straight into your workflow. It makes content curation effortless and dependable.

ZenQuotes API: Tools and Capabilities (4)

These four tools allow you to programmatically check the status of the service, grab a single random quote, get the featured daily quote, or pull large batches of inspirational content.

#	TOOL	DESCRIPTION
01	<code>check_api_status</code>	Checks if the ZenQuotes API service is currently running and available for use.
02	<code>get_zen_quotes_batch</code>	Retrieves a large batch of 50 random inspirational quotes in one call.
03	<code>get_random_zen_quote</code>	Fetches a single, randomly selected inspirational quote for immediate use.
04	<code>get_zen_quote_of_the_day</code>	Pulls the official, curated 'Quote of the Day' from ZenQuotes.

See It in Action

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

U Get a random inspirational quote using ZenQuotes.



I've retrieved a random inspirational quote! It's by [Author]: "[Quote content]". Would you like to see more quotes or the official quote of the day?

U Show me the quote of the day.



I've retrieved the official ZenQuotes Quote of the Day! It is currently identified as: "[Quote content]" by [Author]. I can assist you with a batch of similar quotes if you'd like.

U Get a batch of 50 inspirational quotes.



I've retrieved a batch of 50 inspirational quotes! They feature thoughts from various authors on wisdom, success, and life. I can provide the full list or summarize the main themes for you.

Frequently Asked Questions

01 Can I use ZenQuotes API for quotes from specific authors?

While the tool doesn't filter by author name directly, you can retrieve large batches using ``get_zen_quotes_batch`` and then process the metadata within your AI client to analyze content distribution across different sources.

02 Is there a limit on how many quotes I can get with ZenQuotes API?

The dedicated tool, ``get_zen_quotes_batch``, allows you to retrieve up to 50 inspirational quotes in a single call. You'll need to structure your agent prompts if you require more than that.

03 Do I need an API key to use the ZenQuotes API MCP?

No, because ZenQuotes is a free and open service, this MCP requires no API keys. You just connect it through Vinkius and start using your agent.

04 What if I need to check if the quote source is working before starting?

Use the `check_api_status` tool first. This confirms that the ZenQuotes API service is fully operational, saving you time and frustration later on.

05 Is this MCP suitable for commercial use?







Yes. Since the data comes from a public source and the content is delivered through your agent's workflow, it works perfectly for building revenue-generating applications or marketing tools.

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>"zenquotes-api": { "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

ZenQuotes API 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 ZenQuotes API. 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	ZenQuotes API MCP
Server ID	019d84a2-e95d-7102-bb75-1ba090cab90
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/zenquotes-api.