MCP SERVER | NO CODE | CLOUD HOSTED

Exa MCP

Find Contextual Web Data, Not Just Keywords

Exa MCP finds what you actually need on the web using contextual search. It goes way beyond basic keyword matching, understanding the intent behind your questions to surface high-quality, curated results and extract clean text from multiple sources automatically.

A+ Quality Score 98.33/100

neural-search | semantic-similarity | web-data | contextual-search | content-curation | ai-search

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

SCAN TO VIEW

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.

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.

Exa MCP

10 tools available
Cloud-hosted on Vinkius

This connector lets your AI agent look up information across the entire internet like a highly specialized researcher—but without you having to copy and paste 50 URLs. Instead of relying on simple keyword matching, it uses advanced semantic search to understand the *meaning* behind your request. You can ask complex questions about niche topics or compare concepts between different fields, getting answers based on contextual understanding. It also lets you find articles that are similar in topic, not just keywords, and then pull out clean, readable text from those pages for analysis. By connecting this MCP through Vinkius, your agent gets access to powerful web data discovery right where you're building your workflows. You can use it to analyze complex source material or research an entire domain without ever leaving your workspace.

Core Capabilities

01 — Perform conceptual searches

Run advanced neural searches that understand the meaning and intent of a query, not just matching specific words.

02 — Analyze content similarity

Find web pages or articles that are semantically related to a specific URL you provide for comparative research.

03 — Extract raw text from multiple sources

Fetch cleaned, usable text and metadata simultaneously from several URLs in one single request.

04 — Search within defined domains

Restrict your search to a specific website or knowledge base for focused research or documentation lookups.

05 — Retrieve web data and metadata

Get structured information about search results, helping you identify the most authoritative sources.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and provide your Exa API Key in your client settings.

02 Your AI agent sends a natural language request—for example, 'Find articles similar to X.'

03 The MCP executes the appropriate search or extraction tool, returning structured, clean web content directly back to your agent for use.

The bottom line is you get powerful, context-aware web data retrieval without having to write complex API calls yourself.

Built For

This MCP is built for researchers and content professionals who need deep, contextual insights from the web. If your job involves synthesizing information across multiple disparate sources or proving a thesis using external data, this tool saves hours of manual digging.

Academic Researcher

Uses search_neural to track down recent pre-prints and find semantically similar academic papers for literature reviews.

Content Strategist

Runs find_similar_with_contents against competitor links to gauge the current online conversation around a topic.

Technical Writer

Uses get_contents and search_domain to pull clean documentation text from specific technical sites for manuals.

What Changes When You Connect

01 Stop relying on simple keyword hits. Use search_neural to ask your agent complex conceptual questions and get answers based on the web's deep meaning.

02 Don't waste time copying text from multiple tabs. The get_contents tool pulls clean, readable content and metadata from several URLs in one go.

03 Need to analyze a competitor? find_similar is better than simple searches because it finds pages that are conceptually related, not just keyword-matched.

04 Want to keep your research focused? Use search_domain to limit the scope of your investigation to a single website's documentation or site section.

05 Improve your prompts automatically. The MCP includes tools like 'Autoprompt' to optimize your queries for the highest accuracy before you hit send.

Real-World Applications

Mapping out competitive research gaps

A content strategist wants to see what other blogs are writing about a niche topic. They use find_similar_with_contents on their own best-performing article, and the agent returns five semantically related competitor articles along with the full text for immediate analysis.

Synthesizing complex academic findings

A researcher needs to know about 'room temperature superconductors.' Instead of running ten separate Google searches, they run a neural search and receive a curated set of high-quality sources (pre-prints, university blogs) ready for content extraction.

Drafting technical documentation updates

A developer needs to update an internal guide on a specific API. They use search_domain against the official vendor site to pull all necessary code snippets and parameter definitions, ensuring accuracy before writing anything.

Patterns to Avoid

Treating web data like simple keywords

✗ AVOID

Asking a generic AI agent for 'AI search tools' gets hundreds of basic links that don't address the specific type of tool you need (like an MCP).

✓ INSTEAD

Use search_neural instead. Frame your question conceptually, such as 'What are advanced methods for connecting AI agents to external services?' This forces the search engine to understand context.

Forgetting that web pages aren’t just text

✗ AVOID

Copying content from a PDF or complex website often results in garbled, unreadable text filled with formatting junk.

✓ INSTEAD

Use get_contents. This tool is designed to strip away the noise and deliver clean, usable body text and metadata directly.

The Right Fit

Use this MCP if your core problem involves information discovery or synthesis from external web sources. If you need to understand *context* (e.g., 'What are related concepts?'), use the neural search tools. If you only care about finding exact, literal matches for a single term in one place, then a basic keyword tool will do. But if your workflow involves comparing ideas or extracting clean data from multiple sources, this is your connector. Don't connect this just because it's available; connect it when the web's context-aware power is required.

Diving into research means endless clicking and copy-pasting.

Today, if you need to understand a niche topic or compare three different industry standards, your process looks like this: You open five browser tabs. You run Google searches on each one. Then,

With this MCP, that entire manual loop disappears. Instead of managing multiple tabs and messy clipboard data, you ask your agent one question. It uses semantic search tools to pull in all relevant

for every link, you have to click through the article and manually copy-paste the key paragraph into a single document just so you can start writing.

information—the articles, the technical specs, the related concepts—and gives you clean text ready for analysis.

Exa MCP delivers actionable web context.

The manual steps that vanish include: opening ten different search engines; running fifteen keyword variations; and spending hours cleaning up badly formatted, copy-pasted text from various source types. You no longer have to synthesize the data yourself.

Now, you get a single, coherent stream of high-quality web context, structured by your agent and ready for immediate use in whatever task you're working on.

Exa: 10 Powerful Web Discovery Tools

Use these tools to search the web using semantic understanding, extract content, or analyze related information without manual steps.

#	TOOL	DESCRIPTION
01	`answer`	Gets an AI-generated answer based on web content context.
02	`find_similar`	Finds pages that are semantically similar to a given URL, useful for competitive analysis.
03	`find_similar_with_contents`	Identifies pages similar to a source and includes the actual text content for review.
04	`get_contents`	Extracts clean, readable text and metadata from specified web page IDs or URLs.
05	`search_domain`	Performs a search limited to a specific website's domain, ideal for documentation lookups.
06	`search_keyword`	Executes traditional keyword-based searches across the web.
07	`search_neural`	Conducts conceptual or research-topic searches that understand intent using neural embeddings.
08	`search_recent`	Searches for content published recently, useful for news and trending topics.
09	`search`	Performs a general web search returning titles, URLs, and relevance scores.
10	`search_with_contents`	Searches the web while simultaneously extracting page content for review.

See It in Action

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

U Neural search for 'the latest breakthroughs in room-temperature superconductors'.

I've performed a neural search for you. I found 10 high-quality links including recent pre-prints from arXiv, specialized science blogs, and official university press releases. Which source would you like me to extract the content from?

U Find 5 links similar to 'https://openai.com/blog/instruction-following'.

Fetching similar content... I've identified 5 semantically related articles, including papers on RLHF from Anthropic, Google DeepMind's blog on fine-tuning, and technical breakdowns from HuggingFace. Shall I retrieve the metadata for these?

U Extract the cleaned text content from 'https://arxiv.org/abs/2312.00752'.

Extraction complete! I've retrieved the cleaned text content for the arXiv paper (ID: 2312.00752). It contains the abstract, full body text, and author metadata. Would you like me to summarize the key findings for you?

Frequently Asked Questions

01 **How does Exa MCP search differently than standard Google searches?**

Exa uses neural embeddings to understand the meaning of your query, not just the words. It finds results based on context and intent, making it ideal for conceptual research topics.

02 **Can I use Exa MCP to find content from a specific website?**

Yes. You can restrict searches using search_domain to focus only on one site's documentation or knowledge base, ignoring the rest of the web.

03 **What is the best tool for finding related articles?**

For competitive analysis, use find_similar. If you need both similar pages and their actual text content, use find_similar_with_contents.

04 **Does Exa MCP only give me links, or does it extract the text?**

It does both. You can run get_contents to pull clean text from a URL, and search_with_contents will perform the web search while extracting the page content simultaneously.

05 **Is Exa MCP useful for coding documentation lookups?**

Absolutely. Use search_domain to target specific developer sites or internal knowledge bases, ensuring you only pull relevant technical details.

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 `"exa-alternative": { "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

Exa 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 Exa. 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	Exa MCP
Server ID	`019dd0ec-cc9c-72a3-9cf4-c1cd808fcfea`
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/exa-alternative.