MCP SERVER | NO CODE | CLOUD HOSTED

Elastic Enterprise Search MCP for AI Agents

Manage complex corporate knowledge bases and document indexing workflows

Elastic Enterprise Search MCP connects your AI agents directly to robust enterprise search engines. It lets you manage complex document indexes, run deep contextual queries across multiple data scopes, and audit usage metrics—all through natural conversation. Stop writing boilerplate API calls; start asking questions about your company's knowledge base.

A+ Quality Score 100/100

enterprise-search | indexing | semantic-search | workplace-discovery | data-retrieval | search-analytics

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/elastic-enterprise-search

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.

Elastic Enterprise Search MCP

6 tools available
Cloud-hosted on Vinkius

Need to find something buried deep in your corporate documents? This MCP gives your AI agent full control over your entire search environment. Instead of manually checking Kibana dashboards or running complex REST queries, you simply ask your agent what you need. It manages the connections and executes everything from listing available search engines to performing detailed searches on specific document collections.

The system handles bulk data ingestion for updating indexes and provides precise analytics tracking, so you know exactly which parts of your knowledge base are popular and which ones are failing. When you subscribe through Vinkius, your agent gains access to this powerful capability alongside thousands of others, letting you manage everything from search queries to document lifecycle status without ever leaving your chat interface.

Core Capabilities

01 — List all deployed search engines

Retrieves a list of every available search engine and its current operational status.

02 — Get detailed information on an engine

Fetches comprehensive metadata and configuration details for a single, specified search engine.

03 — Search documents within an engine

Runs natural language queries across the content of one or more engines to find relevant document snippets.

04 — List indexed documents metadata

Retrieves a list and basic details about all documents currently stored in a designated search engine.

05 — Ingest new JSON documents

Sends bulk payloads of newly created data to be processed and stored within specific schemas.

06 — Generate usage analytics reports

Calculates usage insights, tracks click logs, and generates performance metrics for search activity.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/elastic-enterprise-search** — connect your AI agent in three steps.

01 Subscribe to the Elastic Enterprise Search MCP on Vinkius.

02 Provide your specific Elastic Enterprise Search URL and API Key (you can find this key in Kibana's Stack Management).

03 Use your AI client to ask questions. Your agent will interpret your request, execute the necessary search or indexing operations, and return a clean summary of the results.

The bottom line is that you talk naturally, and the MCP handles all the complex back-end API calls needed for enterprise search management.

Built For

This is essential for Search Engineers who need to monitor engine health without manual API calls. It's also perfect for Data Analysts who want to audit usage metrics and spot performance gaps using simple conversation, rather than complex dashboard queries.

Search Engineer

Monitors multiple search engine configurations, tests query relevance, and validates index schemas without writing test code.

Software Developer

Indexes new JSON data into the knowledge base or verifies search results directly from their IDE chat interface.

Data Analyst

Audits click logs and generates usage analytics to identify content gaps or measure feature adoption rates over time.

What Changes When You Connect

01 Stop debugging search relevance manually. Your agent can test query performance across multiple engines using the 'get engine' capability, letting you confirm system health instantly.

02 Update your company knowledge base easily. Use the 'index documents' tool to bulk upload new JSON data without touching a pipeline or writing code.

03 Pinpoint content gaps instantly. The 'analytics' function allows you to audit search usage and see exact click logs, telling you precisely what information employees are looking for but can’t find.

04 Consolidate discovery. You don't need dozens of dashboards; your agent can list all engines and run targeted searches across them using the core 'search' capability.

05 Deepen data visibility. The MCP allows you to inspect index layouts and metadata by calling the engine details tool, giving you a clean view of the underlying data structure.

Real-World Applications

A team needs to audit product search performance

Instead of running weekly reports on click-through rates, an analyst asks their agent: 'Show me the search analytics for e-commerce products.' The agent executes the appropriate tool and delivers a summary showing top queries and overall traffic trends.

A developer needs to onboard new documents

A developer has finished writing 100 new internal guides. They instruct their agent: 'Index these 100 JSON files into the help center.' The agent uses the indexing tool, validating that all data is stored correctly and synchronously.

The operations team needs to verify system health

An ops engineer wants to know if a specific content area is even indexed. They ask their agent: 'List documents for the marketing engine.' The agent replies with a clean list of metadata, confirming the scope and size of the index.

A user needs to find information across departments

A manager asks: 'What's the policy on remote work?' The agent doesn't just search one place; it uses the core search capability to check multiple connected engines simultaneously, providing a comprehensive answer from various sources.

Patterns to Avoid

Over-relying on single document searches

✗ AVOID

A user asks their agent to find 'API integration guides' and the agent only runs a search against one engine, missing relevant documents in another department's index.

✓ INSTEAD

Always confirm all available engines first using the list engines tool. Then, ask your agent to run the core search across multiple specified engines for comprehensive results.

Manually listing document metadata

✗ AVOID

A developer tries to manually track which documents need updating by calling 'list documents' repeatedly and copying raw JSON arrays.

✓ INSTEAD

Use the indexing tool to manage updates in bulk. If you only need a list, use the dedicated list documents capability; don't try to pull records piecemeal.

Ignoring usage patterns

✗ AVOID

An ops team assumes a feature is failing because nobody found it, without checking real metrics.

✓ INSTEAD

Always audit performance first. Run the analytics tool to check click logs and query volume before deciding if content needs fixing or boosting.

The Right Fit

Use this MCP if your core pain point is managing a large, complex knowledge graph where data lives in multiple indexed silos. You need an agent capable of sophisticated search (like the 'search' tool) AND lifecycle management (like 'index documents'). Don't use it if you only have one small database and just need simple CRUD operations; a basic API connector will suffice. If your goal is simply to read raw data without context, stick to a direct database connection instead of an enterprise search layer like this MCP.

Elastic Enterprise Search: Managing Corporate Knowledge Bases via AI Agents

Today, finding specific information means jumping through hoops. You log into the intranet for policy docs, then switch to Jira for project specs, and finally go to a separate knowledge base for technical guides. This constant tab-switching and context-switching is exhausting, and often you only get partial answers.

With this MCP, your agent handles the complexity behind the scenes. Instead of juggling five different dashboards, you simply ask: 'What's the policy on remote work?' The system executes multiple searches across all connected engines, pulling together a single, accurate answer from every relevant source.

How Elastic Enterprise Search Improves Document Indexing with AI Agents

Manual document ingestion is slow. You have to write scripts that read files, validate schemas, and then send them to the search engine in batches. A single failed record can halt the entire process.

Now, you tell your agent: 'Index these 50 new quarterly reports.' The MCP takes care of the bulk payload ingestion, triggering native pipeline mappings automatically. You get verified data stored synchronously without writing a line of indexing code.

6 Tools for Elastic Enterprise Search Indexing and Query Management

Use these tools to list engines, perform complex searches on documents, ingest new data in bulk, and generate detailed performance analytics.

#	TOOL	DESCRIPTION
01	`list_engines`	Retrieves a list of all deployed search engines managed by the system.
02	`get_engine`	Fetches detailed configuration and status information for a single engine.
03	`search`	Finds specific documents across an engine using natural language queries.
04	`list_documents`	Displays the metadata and basic details of all indexed content in a given search engine.
05	`index_documents`	Accepts new JSON data payloads to update or add documents to specific schemas.
06	`analytics`	Calculates and retrieves detailed search usage metrics and performance reports.

See It in Action

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

U What were our top search topics last month?

Here's a breakdown of your activity over the past 30 days:

METRIC	VALUE
Total Queries	8,520
Avg. CTR	7.1%
Top Query	'SaaS integration costs' (450 queries)

We see a big spike in questions about billing structure; maybe we need to update the pricing guide.

U Show me the status of our product documentation engine.

The 'product-docs' engine is currently: **Operational**.

- **Document Count:** 12,450 records
- **Last Indexed:** 2 hours ago (Success)
- **Schema Version:** v3.1

The configuration looks good. Everything seems healthy.

U Can you find any articles about migrating to the new cloud platform?

Searching 'help-center-docs' for 'cloud migration'...

I found 4 relevant documents:
Cloud Migration Guide v2.0* (Score: 18.9) - Full text available.
Migration Checklist* (Score: 15.2) - Metadata only.
API Endpoint Changes* (Score: 12.5) - Partial snippet found.

Frequently Asked Questions

01 **Can my agent list all available search engines in Elastic?**

Yes. Use the 'list_engines' tool. The agent iterates through your engine containers, managing logical indexing schemas and providing a complete map of your search spaces.

02 **How do I index a batch of JSON documents via chat?**

Use the 'index_documents' tool. Provide the engine name and a JSON array of your documents. The agent will command the bulk payload ingestion, triggering native pipeline mappings to store your data synchronously.

03 **Can I check the search analytics for a specific engine through the agent?**

Absolutely. The 'analytics' tool generates precise internal metric tracking for your engine. It will isolate usage insights and calculate click log data, allowing you to monitor search performance natively.

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 `"elastic-enterprise-search": { "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

Elastic Enterprise Search 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 Elastic Enterprise Search. 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	Elastic Enterprise Search MCP
Server ID	`019d758e-9110-73c4-9ebc-a669f1a0382d`
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/elastic-enterprise-search.