MCP SERVER | NO CODE | CLOUD HOSTED

Ada MCP for AI Agents

Automating Customer Support and Knowledge Base Management

Ada connects your AI agent directly to advanced customer service automation tools. Your agent can monitor conversation history, sync user metadata, and manage your entire knowledge base through natural language commands.

A+ Quality Score 100/100

conversational-ai | customer-service-automation | chatbot | knowledge-base | support-automation | user-metadata

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/ada

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.

Ada MCP

4 tools available
Cloud-hosted on Vinkius

Connect this MCP to give your AI agent deep visibility into your customer support operations. Instead of having your team manually jump between a CRM, an internal wiki, and a chat log just to answer one question, your agent handles the orchestration for you. It can pull up user profiles, check the status of past conversations, and even read articles from your knowledge base to give accurate answers in real time.

Need to update what the bot knows? You can use the MCP to create new help articles directly into Ada's system, immediately improving the quality of responses. Because all these capabilities are managed through a single connection point, you get comprehensive support automation without needing specialized developer endpoints. Just connect your preferred AI client via Vinkius and start managing your entire customer service ecosystem using natural conversation.

Core Capabilities

01 — Track conversation history

List active and past support chats handled by the Ada bot.

02 — Fetch user profiles

Retrieve profile information and custom data points for a specific Ada end user.

03 — List knowledge articles

Retrieve the catalog of help articles used by the Ada AI agent to answer customer queries.

04 — Create new knowledge articles

Add a new text article to the Ada knowledge base, which immediately improves bot responses.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP on Vinkius.

02 Enter your Ada Platform Token and Handle credentials.

03 Your AI client accesses all support tools through natural language conversation, letting you manage the entire knowledge base.

The bottom line is, you connect once and get instant access to all of Ada's customer service data streams.

Built For

This MCP helps Support Operations Teams who are tired of logging into multiple dashboards just to audit conversation quality. It's for Customer Experience Managers needing real-time oversight and Developers building custom chat applications that need structured data.

Support Operations Manager

Needs to review patterns in support conversations, check if users have synced metadata, and quickly update the internal knowledge base.

Customer Experience Analyst

Monitors conversation transcripts to identify gaps in automation or common failure points so they know where to add new articles.

Technical Developer

Integrates Ada's conversational AI into custom apps, using the MCP to fetch user data and article content programmatically.

What Changes When You Connect

01 Review support trends instantly. Use `list_conversations` to pull together activity data from both automated resolutions and human handoffs.

02 Enrich context for agents. The `get_end_user` tool lets you fetch user metadata, ensuring the AI agent knows who it's talking to before responding.

03 Maintain up-to-date content. With `list_articles`, you get a clear view of your current help documentation catalog.

04 Improve bot accuracy on demand. Use `create_article` to publish new knowledge base entries and raise the quality of automated answers immediately.

05 Simplify compliance tasks. You can monitor conversations and manage data privacy requests through one centralized chat interface.

Real-World Applications

Identifying a failing feature

A support operations manager uses the MCP to run `list_conversations` for the past week. They notice that 60% of chats involving 'billing' are being escalated to human agents, indicating a critical knowledge gap.

Onboarding a new product line

A CX manager uses `list_articles` to confirm the current documentation. Then, they use `create_article` to add three brand-new guides about the product's latest features, ensuring the bot is ready for launch day.

Resolving complex user queries

A developer needs context on a specific customer. They ask their agent to use `get_end_user` with an ID, which fetches account status and custom metavariables needed to resolve the chat.

Auditing bot performance

An AI training lead needs to audit if a specific policy was covered. They check `list_articles` first, then use `get_end_user` to verify which version of the user profile was used during a key conversation.

Patterns to Avoid

Assuming all data is in one place

✗ AVOID

Trying to figure out if a new topic needs documentation by manually checking multiple systems, or asking the AI agent to 'just know' the answer.

✓ INSTEAD

Always use `list_articles` first to check your existing content catalog. If the answer isn't there, then use `create_article` to formally add it.

Ignoring user context

✗ AVOID

The agent gives a generic answer because it doesn't know if you are a premium or free tier customer.

✓ INSTEAD

Before asking any support question, run `get_end_user` to fetch the necessary account metadata. This provides crucial context for accurate responses.

Treating conversations as static logs

✗ AVOID

Only looking at conversation transcripts without understanding what happened or why.

✓ INSTEAD

Use `list_conversations` to pull the history, and then cross-reference that data with user details using `get_end_user` for a complete picture.

The Right Fit

Use this MCP if your primary challenge is integrating multiple sources of customer knowledge—conversation logs, user profiles, and documentation—into one place. You need an AI agent that can synthesize answers from disparate data points. However, don't use it if you simply want to send bulk internal emails or manage ticketing system assignments; those are communication-focused tools. If your only goal is generating reports on conversations without accessing the underlying article content, a dedicated reporting tool will be better. This MCP excels at enrichment and knowledge management.

Ada MCP for AI Agents: Managing Support Knowledge Articles

Right now, updating your chatbot's answers is a manual headache. You have to write the article in Confluence, then copy it into the support portal, and finally ask someone on the engineering team to sync it all up. It’s slow, error-prone, and often outdated by the time it goes live.

With this MCP, you just need to tell your agent what the new information is. You run `create_article` and instantly publish that content to Ada's knowledge base. The bot sees it immediately. That takes the bottleneck out of your documentation workflow.

Ada MCP for AI Agents: Syncing User Data and Conversation History

Before, if a user called in, you'd have to ask them for their account ID, then open the CRM, search by ID, copy over basic metadata like 'premium status,' and finally paste that into your notes. It’s three systems and ten clicks.

Now, simply asking your agent to run `get_end_user` pulls all that profile data automatically. The AI client provides you with a clean JSON object containing the user's full context right when you need it.

4 Ada Tools for Conversational AI Knowledge Management

Use these four tools to list conversation history, retrieve user metadata, catalog articles, and create new content.

#	TOOL	DESCRIPTION
01	`list_conversations`	Retrieves a list of both ongoing and past support chats handled by the Ada bot.
02	`get_end_user`	Fetches detailed profile information and custom data points for a specific user ID within Ada.
03	`list_articles`	Retrieves the full catalog of help articles used by the AI agent to answer customer questions.
04	`create_article`	Adds a new text article to your knowledge base, which immediately helps improve the bot's answers.

See It in Action

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

U Show me the last 5 support chats and if there are any trends in billing questions?

Last 5 Conversations Summary

ID	STATUS	TOPIC
C-901	Resolved (Auto)	Billing Inquiry
C-887	Handoff to Human	Feature Request
C-852	Resolved (Human)	Login Issue

Analysis: We see a spike in billing inquiries this week. 3 of the last 5 conversations mention 'overcharge.' This suggests we need new documentation on our pricing tiers.

U I have user ID 4492. What's their status and any custom details I should know?

User Profile: John Doe (ID: 4492)

- **Account Status:** Premium Tier
- **Last Interaction:** Yesterday, support chat.
- **Custom Metadata:** `onboarding_date`: 2023-11-01; `department`: Marketing. This user is a key contact and should be treated with high priority.

U What articles do we have on setting up the API endpoint?

Ada Knowledge Base Article Catalog

We currently have 14 published articles covering technical topics. Here are a few relevant ones:
• Getting Started: Basic Setup (Article ID: KB-001)
API Endpoint Configuration Guide (Article ID: KB-022) - This is the guide you're looking for.*
• Error Code Dictionary.

Frequently Asked Questions

01 **How does the Ada MCP help me keep my chatbot knowledgeable?**

It gives you direct control over your knowledge base. If a support agent learns something new, you can use the MCP to create and publish a brand-new article immediately, making sure the bot has accurate information right away.

02 **Can I get customer data into my AI client using this Ada MCP?**

Yes. You can fetch detailed user profiles for any specific ID. This includes core account details and custom metadata, letting your agent give highly personalized advice.

03 **What kind of support trends can I see with the Ada MCP?**

You can retrieve conversation history to spot patterns. By listing conversations, you can analyze which topics are generating the most chats or require human handoff, guiding where your team needs to focus.

04 **Is this better than just reading the Ada dashboard?**

It's faster and more flexible. Instead of logging into a separate dashboard, you ask your agent via your AI client. The MCP handles all the data retrieval and presents it in a conversation format.

05 **What if I need to update an article but don’t know where?**

First, use the listing function to view the existing article catalog. This shows you everything that’s already published before you decide which topics need updating or adding.

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 `"ada": { "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

Ada 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 Ada. 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	Ada MCP
Server ID	`019d7546-93b2-7010-88fa-0f79b648a326`
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/ada.