MCP SERVER | NO CODE | CLOUD HOSTED

Doppler

A+ Quality Score 100/100

secrets-management | environment-variables | doppler | api-keys | config-management

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026

vinkius.com/mcp/doppler

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.

Doppler MCP

12 tools available
Cloud-hosted on Vinkius

One Click on Vinkius — From Prompt to Execution

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

Built For

12 Tools Available

#	TOOL	DESCRIPTION
01	change_secrets	Provide project_slug, config_name and a JSON object mapping secret names to values. For example: {"DATABASE_URL":"postgres://...","API_KEY":"sk-..."}. Existing secrets not included are not modified. Add or update secrets in a Doppler config
02	delete_secrets	Provide project_slug, config_name and comma-separated secret names. WARNING: deleted secrets cannot be recovered. If a secret inherits a value from a parent, it reverts to that value. Delete secrets from a Doppler config
03	get_config	Returns config name, project, root status, associated environment template, creation date and locked status. Get details for a specific Doppler config
04	get_account	Returns account email, name, and token metadata (type, scope, permissions). Use this to verify your token is working correctly and understand its access level. Get the current Doppler account details
05	get_project	Provide the project slug (e.g. "my-api-project") and optionally the workspace slug. Get details for a specific Doppler project
06	get_secret	Returns the secret name and its resolved value with fallbacks from parent environments applied. Get a specific secret value from a Doppler config
07	list_activity_logs	Each entry shows who performed what action, when and the affected config. Optionally filter by config_name. Useful for security auditing and compliance. List activity logs for a Doppler project
08	list_configs	Each config represents a deployment environment (development, staging, production) and contains its own set of secrets. Returns config name, project slug, root status and environment template used. List configs (environments) for a Doppler project
09	list_environments	g. development, staging, production, preview). Returns environment name, slug and whether it is the default environment. List Doppler environment types
10	list_projects	Optionally filter by workspace slug. Each project contains configs (environments) and secrets. Returns project name, slug, description, and creation date. List Doppler projects
11	list_secrets	Returns each secret's name, computed value (with environment fallbacks applied), visibility status. Provide the project_slug and config_name. List all secrets for a Doppler config
12	list_workspaces	A workspace is the top-level organizational unit in Doppler that groups projects. Returns workspace name, slug and creation date. List all Doppler

#	TOOL	DESCRIPTION
		workspaces

See It in Action

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

U Show me all configs for my 'backend-api' project.

I found 3 configs for 'backend-api': 'dev' (development, root config), 'staging' (staging, inherits from dev) and 'prod' (production, inherits from staging). Each has its own set of secrets.

U Update the DATABASE_URL secret in my prod config to point to the new database.

Done! I've updated the `DATABASE_URL` secret in the prod config of your backend-api project. The new value has been applied and will be available to all services using this environment.

U Who changed secrets in my project in the last week?

I found 14 activity events in the last 7 days: 3 secret updates by alice@company.com, 8 secret reads by the CI/CD service account, 2 config changes by bob@company.com and 1 new secret added by admin@company.com.

Frequently Asked Questions

01 **How do I create a Doppler Service Token?**

Log in to the [**Doppler Dashboard**](https://dashboard.doppler.com), select your project, go to **Settings** > **Tokens** and click **Generate Token**. Choose the scope (project + config/environment), set the access level (Read or Read+Write) and copy the token immediately — it won't be shown again.

02 **Can I update multiple secrets at once?**

Yes! Use the `change_secrets` tool with a JSON object mapping names to values, e.g. `{"DATABASE_URL":"postgres://new-host","API_KEY":"sk-new"}`. This creates or updates all specified secrets in a single atomic operation.

03 **What is the difference between a Personal Token and a Service Token?**

A **Personal Token** is scoped to your user account and can access all workspaces and projects you have permission for. A **Service Token** is scoped to a specific project and config, with either read-only or read+write access. Service tokens are recommended for CI/CD and automated integrations, while personal tokens are better for development and admin tasks.

04 **Can I view the activity history for a project?**

Yes! Use the `list_activity_logs` tool with the project_slug to see all audit events (secret reads, writes, config changes, user additions). Optionally filter by config_name to see activity for a specific environment only. Each log entry shows who performed the action, when, and what was affected.

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 `"doppler": { "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

Doppler 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 Doppler. 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	Doppler MCP
Server ID	`019d842f-e053-713d-89fb-280d179b7cfc`
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/doppler.