MCP SERVER | NO CODE | CLOUD HOSTED

Sentry MCP for AI Agents

Investigate production errors and stack traces in real time.

Sentry MCP gives your AI agent direct access to Sentry's application performance monitoring. It lets your agent pull raw exceptions, read stack traces, and manage issue statuses in real time. Use it to triage production crashes and resolve bugs directly from your AI client.

A+ Quality Score 100/100

error-tracking | performance-monitoring | debugging | stack-trace | incident-management

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JULY 2026
vinkius.com/mcp/sentry

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.

Sentry MCP

10 tools available
Cloud-hosted on Vinkius

The Sentry MCP gives your AI agent direct access to Sentry's application performance monitoring. You spend way too much time switching tabs between your code and the Sentry dashboard just to figure out why a production build is melting down. This MCP changes that by letting your AI agent do the heavy lifting for you. Instead of hunting for a specific error ID, you can just tell your agent to find the latest issues in your project and summarize what's actually happening. It pulls the stack traces directly into your chat window so you can see the exact line of code that failed without leaving your workspace. It's a huge time saver when you're trying to catch a bug before it hits your users. You can even have your agent handle the boring stuff, like closing out issues once you've pushed a fix. This is part of the Vinkius catalog, which makes it easy to plug into your existing workflow with just a few credentials. It gives you a direct line into your operational data so you can stop guessing and start fixing.

Core Capabilities

01 — Fetch latest crash logs

The agent pulls the most recent errors from your project to show you what is breaking right now.

02 — Analyze stack traces

Your AI can read the full technical trace of a crash to pinpoint the exact line of code at fault.

03 — Close resolved bugs

You can tell your agent to mark issues as resolved without having to manually update the dashboard.

04 — Scan project structures

The agent can list all projects and organizations to help you navigate different environments.

05 — Audit team memberships

You can check who has access to your Sentry workspace or see how teams are organized.

06 — Monitor live errors

The agent can pull specific event details to help you see the context behind a single failure.

One Click on Vinkius — From Prompt to Execution

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

01 Connect the Sentry MCP to your Vinkius dashboard.

02 Provide your Organization Slug and a scoped Auth Token.

03 Ask your agent to summarize recent crashes or find specific error details.

The bottom line is you get instant access to Sentry's error data without leaving your IDE.

Built For

The Sentry MCP is built for the engineers who live in the terminal and the ops teams who need to stop clicking through nested dashboards to find a single failing line of code.

On-call Engineer

They use this on Tuesday nights to quickly identify which production service is throwing the most errors and why.

Frontend Developer

They use this to pull specific crash traces into their chat window to see exactly what state caused a UI component to fail.

Technical Founder

They use this to ask for a high-level summary of critical bugs across all projects before a daily standup.

What Changes When You Connect

01 You can triage production crashes faster by using `list_issues` to see what is breaking immediately.

02 You get deep context on bugs because `get_issue_details` feeds the exact stack trace to your agent.

03 You keep your dashboard clean by using `resolve_issue` to mark fixed bugs without opening the web UI.

04 You can audit your workspace security by using `list_organization_users` to see who has access.

05 You can manage multiple apps easily by using `list_projects` to switch contexts for different services.

06 You get a clear history of app behavior by using `list_events` to see the sequence of actions leading to a crash.

Real-World Applications

Production Crash Triage

A deployment causes a spike in errors. An engineer asks the agent to find the most frequent crash and summarize the stack trace using `list_issues` and `get_issue_details`.

Database Deadlock Investigation

A specific error ID is flagged. The user asks the agent to fetch all metadata for that issue to see if it correlates with a recent database migration.

Workspace Security Audit

A founder wants to know who has access to the production environment. They ask the agent to list all users and teams in the organization.

Daily Standup Summary

Before a meeting, a lead asks the agent to list all high-priority issues across all projects to get a quick status update on the day's bugs.

Patterns to Avoid

Using it as a notification tool

✗ AVOID

Asking the agent to ping you every time an error occurs.

✓ INSTEAD

Use a dedicated alerting tool for pings. Use this MCP to investigate the details of those alerts once they arrive using `get_issue_details`.

Irreversible deletions

✗ AVOID

Using `delete_issue` to clear out a bug you haven't actually finished fixing yet.

✓ INSTEAD

Use `resolve_issue` to mark a bug as finished. Only use `delete_issue` for permanent cleanup of data you no longer need.

Web browsing requests

✗ AVOID

Asking the agent to search Google for a Sentry error.

✓ INSTEAD

The agent can only see your Sentry data. Ask it to analyze the specific stack trace it retrieves from `get_issue_details` instead.

The Right Fit

Use this if you need your AI to actually read stack traces or list issues in real time. It's the right choice for engineers who want to automate the initial triage of a production crash. Don't use it if you just want a simple notification when something breaks. For that, a dedicated alerting tool like Slack or PagerDuty is better because this MCP is for deep investigation, not simple pings.

Sentry Error Tracking for Faster Debugging

You're in the middle of a deployment and the alerts start screaming. You have to open a new tab, log into Sentry, filter by the right project, find the most recent crash, and then copy-paste a massive stack trace into your notes or your chat tool.

This MCP lets you stay where you're already working. Just ask your agent to find the latest errors and it'll pull the specific stack trace and metadata right into your conversation. You get the answers you need without the tab-switching fatigue.

Sentry Incident Management for DevOps Teams

You can skip the manual work of checking team permissions or listing projects across different

It turns your AI into a proactive incident responder. You stop hunting for data and start acting on it

environments. The AI can scan your organization structure and tell you exactly what's happening across your entire silo.

immediately.

Sentry 10-Tool Production Error Tracking

Use these tools to fetch logs, analyze stack traces, manage issue statuses, and audit your Sentry organization structure.

#	TOOL	DESCRIPTION
01	delete_issue	The delete_issue tool permanently removes an issue from your project. Use it only for cleanup because this action is irreversible.
02	list_events	The list_events tool lists recent events for a specific project. It helps you see the chronological flow of what's happening in your app.
03	get_event_details	The get_event_details tool pulls the full details for a single event. It's the best way to see the specific metadata behind a single occurrence.
04	list_organizations	The list_organizations tool shows all the organizations you have access to. This helps you navigate between different workspace environments.
05	list_projects	The list_projects tool gets a list of all projects within an organization. Use this to scope your AI's search to a specific app or service.
06	resolve_issue	The resolve_issue tool marks an issue as resolved in Sentry. This helps keep your dashboard clean and signals that a bug is fixed.
07	list_organization_teams	The list_organization_teams tool lists all teams associated with an organization. This is useful for auditing permissions or looking up team structures.
08	list_organization_users	The list_organization_users tool lists all users in an organization. Use this to check who has access to your Sentry environment.
09	get_issue_details	The get_issue_details tool pulls the full details and stack trace for a specific issue. This is the primary way to let your agent analyze a crash.
10	list_issues	The list_issues tool lists all errors and issues in a project. Use this to find what's currently breaking in your production environment.

See It in Action

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

U Why did my app crash in the last hour?

I've checked your recent logs. Your app experienced 12 crashes in the last 60 minutes.

Primary Issue:
- **Error:** `TypeError: Cannot read properties of undefined (reading 'auth_token')`
- **Frequency:** 8 occurrences
- **Impact:** Affecting the login flow.

Would you like me to pull the full stack trace for this specific error?

U Get details on issue 6B3VX4921.

Here are the details for issue **6B3VX4921**:

- **Type:** PostgreSQL Deadlock
- **Location:** `payment_processor.ts` at Line 132
- **Context:** This occurred during a heavy write operation.
- **Impact:** 300+ users affected since Tuesday.

It looks like this is tied to the recent migration you pushed.

U Close the bug for the login page.

I've successfully updated the status of the login page error in Sentry. It is now marked as **RESOLVED**. I've also updated your dashboard to reflect this change.

Frequently Asked Questions

01 **How does the Sentry MCP help me debug faster?**

It allows your AI agent to pull real-time error logs and stack traces directly into your chat. You can ask questions about specific crashes and get immediate answers without leaving your workspace.

02 **Can I use the Sentry MCP to manage my issues?**

Yes. You can tell your AI agent to mark issues as resolved or find specific ones in your project. It handles the manual clicking for you.

03 **Does the Sentry MCP work with my current IDE?**

Yes. Since it's an MCP, it works with any compatible client like Cursor, Windsurf, or VS Code, giving you Sentry data right where you write code.

04 **Is my Sentry data secure with this MCP?**

The MCP uses your own scoped Auth Token and Organization Slug. You control exactly what parts of your Sentry environment the AI can see.

05 **Can the AI agent see my whole Sentry account?**

It depends on the permissions of the Auth Token you provide. You can scope it to specific projects or organizations to keep it limited.

06 **What happens when I ask about a crash?**

The agent will search your project for recent errors, find the most relevant ones, and summarize the technical details and stack traces for you.

07 **Can this AI integration actually mark errors as fixed?**

Yes. This agent component possesses mutable write access. If you invoke the prompt properly, it will fire the `resolve_issue` tool, marking the corresponding exception ID completely dealt with inside the Sentry ecosystem. It can also erase bugs fully via `delete_issue`.

08 **What is the difference between inspecting an 'Issue' and an 'Event'?**

An 'Issue' gathers underlying multiple occurrences of the identical stack exception into one overarching master group. In contrast, querying an 'Event' (`get_event_details`) focuses the AI on a strictly singular, point-in-time incidence where the system crashed.

09 Do I need to supply the Organization Slug with every command?

You configure the overarching Organization Slug strictly at startup globally. For project-level filters, let your LLM query `list_projects` first to fetch internal slugs naturally into its memory buffer without constant repeated user inputs.

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 `"sentry": { "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

Sentry 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 Sentry. 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	July 2026
MCP Server	Sentry MCP
Server ID	`019d7606-15e8-705a-b07f-8fa255b79b92`
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/sentry.