MCP SERVER | NO CODE | CLOUD HOSTED

CodeRabbit MCP for AI Agents

Manage Code Review Governance and Team Productivity Metrics

CodeRabbit manages your entire code review process through natural conversation. This MCP lets you control user roles, assign and unassign seats instantly, track detailed PR metrics, and audit every administrative action from any AI agent.

F Quality Score 3.6/100

ai-code-review | pull-request-automation | code-quality | security-analysis | team-productivity | dev-workflow

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

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.

CodeRabbit MCP

9 tools available
Cloud-hosted on Vinkius

Managing an engineering team's code reviews used to mean jumping between dashboards, running reports, and manually updating access lists—a total time suck. Now, connect your CodeRabbit organization to Vinkius and let your AI agent handle the governance part through natural conversation.

Instead of digging into complex menus, you simply ask your agent what's going on with team seats or who needs training. You can instantly list every user, assign bulk seats across 500 members, or even promote an engineer to admin status without ever touching a settings panel. This MCP gives your AI client full control over the review lifecycle—from tracking average complexity scores on pull requests to generating tamper-resistant audit logs for compliance reporting. It turns tedious governance tasks into simple conversation prompts.

Core Capabilities

01 — List all organization members

You can check who is in the organization and filter that list by their role or whether they have a seat assigned.

02 — Assign CodeRabbit seats to users

This function lets you assign active code review seats to up to 500 user IDs per request.

03 — Retrieve organization audit logs

It pulls tamper-resistant records detailing all administrative actions taken within the CodeRabbit environment, useful for compliance.

04 — Demote users from admin to member role

You can safely demote specific users, removing their elevated administrator privileges while keeping them as standard members.

05 — Get the current seat assignment mode

This checks and reports the operational policy for how CodeRabbit seats are currently being assigned in the organization.

06 — Retrieve PR review metrics for a date range

The agent gathers detailed data on merged pull requests, including complexity scores and average time taken for reviews.

07 — Promote users to admin role

This function elevates specified members into the administrator role, granting them full control over settings.

08 — Remove CodeRabbit seats from users

You can unassign seats from users without deleting their entire accounts, which is useful for temporary access restrictions.

09 —

Allows you to change the overall seat assignment policy mode across the organization (requires Enterprise plan).

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP on Vinkius.

02 Provide your CodeRabbit API Key in the organization settings.

03 Use your AI agent to issue commands like 'list all users with unassigned seats' directly through your preferred client.

The bottom line is, you tell your AI client what governance task needs doing, and it handles the connection and execution using CodeRabbit's API.

Built For

This MCP is built for engineering managers, platform engineers, and compliance officers. If your job involves managing who has access to code reviews, tracking team output, or proving audit trails, this is what you need.

Engineering Manager

You use it to get instant visibility into team velocity and quality metrics by querying pull request review data.

Platform Engineer

You automate seat management, role changes, and configuration updates across massive developer organizations.

Compliance Officer

You query the full audit logs to generate reports on administrative actions without needing direct dashboard access.

What Changes When You Connect

01 Get instant visibility into team code review velocity. Instead of downloading reports, you simply ask your agent to retrieve PR review metrics using get_metrics.

02 Maintain strict compliance records effortlessly. The get_audit_logs tool provides tamper-resistant access to all admin actions for SIEM integration.

03 Manage large teams at scale. Use assign_seats or unassign_seats to control hundreds of user seats in bulk, saving hours of manual clicking.

04 Enforce proper team structure with precision. You can promote users using promote_users or demote them using demote_users when roles change.

05 Know your current policy instantly. Use get_seat_mode and update_seat_mode to manage how seats are assigned across the organization.

Real-World Applications

Identifying team members who lack review access

An engineer needs to know which new hires haven't been given code review seats yet. The agent runs list_users, filters by seat assignment status, and provides a clean list of people needing immediate attention.

Auditing an admin change for compliance

A security officer needs proof that the team lead was only given admin rights on June 1st. They query get_audit_logs, filtering by date and action type to generate a perfect report.

Analyzing team efficiency over time

The manager wants to know if the average review time increased after adopting a new codebase. The agent calls get_metrics for the last quarter, providing complex score and time trend data.

Patterns to Avoid

Manually checking every user status

✗ AVOID

The manager opens 15 different tabs in the dashboard to check if specific users have seats assigned. This takes 20 minutes and is prone to human error.

✓ INSTEAD

Instead, ask your agent to list_users and filter by seat assignment status. It runs the query instantly and gives you a summary of who needs attention.

Forgetting role changes

✗ AVOID

A developer leaves the team and their admin privileges aren't revoked, creating an unnecessary security risk.

✓ INSTEAD

Use demote_users to immediately revoke elevated permissions for departing staff. This ensures the user loses administrative control instantly.

Assuming default seat policies are active

✗ AVOID

The team starts assigning seats inconsistently, leading to confusion about who is supposed to have access.

✓ INSTEAD

Use get_seat_mode and update_seat_mode. This forces the agent to check the current policy and allows you to enforce a standardized assignment mode.

The Right Fit

You should use this MCP if your governance concerns center on who can review code, what their role is, or how much effort the team puts into reviews. If you need to track specific PR metrics like complexity scores or average review time, this is your tool. Don't use it if you primarily need to manage user identities outside of CodeRabbit (like HR data), or if you just want a general chat interface for coding help; those require different agent integrations. Use list_users and get_audit_logs when governance and security are the core concerns.

CodeRabbit MCP: Streamlining Code Review Governance with AI Agents

Right now, managing access to code reviews is a nightmare of clicks. You have to jump into the dashboard, run multiple reports, and manually cross-reference who has an active seat versus who needs elevated permissions. This process eats up time that should be spent coding.

With this MCP, you just talk to your agent. Say, 'Who are the three people promoted last month?' Your agent uses its tools to check roles, list users, and give you a precise answer immediately. You get governance control through conversation.

CodeRabbit MCP: Tracking Code Review Productivity Metrics via AI Agents

To assess team performance, developers usually have to pull raw data on every merged PR, then manually calculate the average complexity score or review time across different branches. It's a massive spreadsheet effort.

Now, you ask your agent for 'Q3 review metrics.' It runs get_metrics and returns clean, actionable figures right away. You don't just track activity; you measure actual quality.

CodeRabbit: 9 Tools for Code Review Governance and Metrics

These tools give your agent granular control over user access, role assignments, seat management, and performance data within your code review system.

#	TOOL	DESCRIPTION
01	assign_seats	Assign CodeRabbit seats to a group of users, supporting up to 500 IDs in one request.
02	get_audit_logs	Retrieves the organization's tamper-proof audit logs for compliance reporting purposes.
03	demote_users	Demotes specified users from an admin role back to a standard member status.
04	get_seat_mode	Checks and returns the current system mode for assigning CodeRabbit seats.
05	list_users	Lists all users in the organization, allowing optional filtering by role or seat assignment status.
06	get_metrics	Retrieves key PR review metrics for a specified date range to analyze productivity.
07	promote_users	Promotes specified users, elevating them to the administrator role within CodeRabbit.
08	unassign_seats	Removes existing CodeRabbit seats from users without deleting their underlying accounts.
09	update_seat_mode	Changes the overall seat assignment policy, requiring an Enterprise plan.

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 team members who do not have a CodeRabbit seat assigned or are currently demoted.

I found 3 users that need attention:

- **sarah.chen**: Status: Unassigned Seat (Action needed)
- **marcus.dev**: Role: Member; Seats: No (Needs promotion/seat)
- **new.hire**: Status: Unassigned Seat (Requires initial assignment)

Would you like me to run assign_seats on all of them?

U What were our average PR complexity scores for the last month?

Here are the review metrics for June 2026:

- **Total PRs reviewed:** 185
- **Average Complexity Score (last 30 days):** 7.9/10
- **Median Review Time:** 4.5 minutes
- **Critical Issues Found:** 25

The score is slightly up; should we investigate the repository metrics?

U I need to confirm who has admin rights and what changes were made last week.

Here are the audit logs for the past 7 days:

1. **role.promoted**: The ops team promoted john.doe to Admin (June 25).
2. **seat.assigned**: Three new seats were assigned by admin@company.com (June 28).
3. **config.updated**: Seat assignment mode changed from Manual to Automatic (June 30).

All entries are logged and ready for compliance review.

Frequently Asked Questions

01 **How does CodeRabbit help me manage roles and user access?**

This MCP lets you control who has what permissions through conversation. You can list users, promote members to admin, or demote them instantly without navigating complex settings.

02 **Can I track team productivity metrics using CodeRabbit with AI agents?**

Yes, you can retrieve detailed PR review metrics like average complexity scores and total review times for any date range. This gives you a clear view of the team's actual output quality.

03 **Is CodeRabbit good for compliance auditing?**

Absolutely. It provides access to tamper-resistant audit logs, letting you query every administrative action taken on the organization, which is critical for SIEM reporting.

04 **Do I need to manually assign seats when new employees join?**

No. You can use this MCP to check who has unassigned seats and then bulk assign them across your entire team using a single command prompt.

05 **What if I want to change the overall seat assignment policy?**

You can view the current policy with get_seat_mode, and if needed (and on the Enterprise plan), you can update the mode using update_seat_mode.

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 `"coderabbit": { "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

CodeRabbit 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 CodeRabbit. 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	CodeRabbit MCP
Server ID	`019d7576-8232-7107-9acd-e11bae6e81d6`
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/coderabbit.

© 2026 Vinkius · vinkius.com · support@vinkius.com