MCP SERVER | NO CODE | CLOUD HOSTED

Postmark MCP

Check email stats, track bounces, and send emails.

Postmark MCP manages transactional email delivery directly through your AI client. Use this connector to send high-volume emails in batches or individually, manage all templates, and track real-time bounce rates without leaving your chat interface. It lets you audit sender domains, check performance statistics, and monitor account health programmatically.

A+ Quality Score 100/100

transactional-email | email-delivery | template-management | bounce-tracking | api-integration

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/postmark-extended

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.

Postmark MCP

11 tools available
Cloud-hosted on Vinkius

Connecting Postmark gives your agent full control over transactional email operations. You stop logging into dashboards to check send metrics or look up failed messages. Instead, you just ask your AI client for what you need: 'Show me all bounces from last week' or 'List the current marketing templates.' The MCP lets you manage template content and verify domains used for sending. It handles everything from single-message sends to large batches of emails, ensuring consistent messaging every time. Because this entire catalog lives within Vinkius, connecting your agent means gaining instant access to all these vital email tools in one place.

Core Capabilities

01 — Send transactional messages

Your agent sends single or bulk emails using verified signatures and defined templates.

02 — Monitor delivery performance

You retrieve detailed statistics on email delivery, including sent volume and open rates.

03 — Track bounces and failures

The MCP pulls a history of bounced emails, allowing you to identify failure patterns.

04 — Manage templates and domains

You can query and get details on existing email templates or list all verified sending domains for audit purposes.

05 — Audit account settings

Your agent provides information about your entire Postmark server configuration, including listing available servers.

One Click on Vinkius — From Prompt to Execution

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

01 First, you subscribe to this MCP and provide your unique Postmark Server Token (and optional Account Token) in the Vinkius interface.

02 Next, you trigger a request from your AI client—for example, asking it to 'list recent bounces' or 'send an email.'

03 Finally, the MCP executes the command using your credentials and returns structured data directly back to your conversation.

The bottom line is that you manage complex email infrastructure through simple conversational prompts.

Built For

This connector is for anyone responsible for high-volume, reliable communication. If you're the person who spends too much time clicking between dashboards just to check if an email campaign went out correctly, this MCP saves your sanity.

Marketing Operations Manager

You use it to audit bounce lists and ensure that every automated welcome sequence sends through the correct, up-to-date template.

Software Engineer

You connect your agent to send critical transactional emails—like password resets or payment confirmations—and check delivery stats programmatically before deploying code.

Customer Success Lead

You use it to verify account settings and pull out historical bounce information to troubleshoot complex client communication issues.

What Changes When You Connect

01 You stop manually compiling reports. Need to know how many emails bounced? Just ask your agent; it uses the `list_bounces` tool to pull a real-time history for you.

02 Forget logging into multiple dashboards. You can manage everything—from checking domain status with `list_domains` to reviewing template content using `get_template` —all in one chat window.

03 When sending emails, you don't need separate code snippets. Use the dedicated `send_batch_email` tool to handle large campaigns reliably and quickly.

04 Debugging delivery issues becomes immediate. If a campaign is underperforming, your agent can run both `get_delivery_stats` and `get_outbound_stats` to pinpoint exactly where the failure happened.

05 You gain full control over your infrastructure. You can check the health of your entire setup by running `list_account_servers` or getting granular details using `get_server_info`.

Real-World Applications

The welcome sequence failed to send.

A new user reports they never got the onboarding email. You ask your agent, and it uses `list_outbound_messages` to confirm if an attempt was made, then checks `get_template` to verify the content used. It quickly determines if the problem is sending or template-related.

We suspect a bad sender domain.

Marketing runs a big campaign and sees poor performance. You ask your agent to run `list_domains` and compare that list against recent bounce data from `list_bounces`. This helps isolate if the issue is with an unverified sending source.

We need to send 50 personalized emails quickly.

Instead of writing a loop, you simply ask your agent to execute `send_batch_email` using the appropriate template ID. The system sends all messages at once and provides confirmation metrics.

Audit our entire email setup for compliance.

You tell your agent to run a full audit: check `list_account_servers`, verify templates with `list_email_templates`, and finally pull current performance numbers using `get_delivery_stats`.

Patterns to Avoid

Assuming all emails are sent.

✗ AVOID

A user assumes that because they ran a send command, the email went through. They don't check if any addresses were invalid or if there was an outage.

✓ INSTEAD

Always verify the outcome by running `list_bounces` after sending, and confirm overall health using `get_delivery_stats`. This shows you exactly which messages failed.

Hardcoding templates into code.

✗ AVOID

A developer writes a script that manually inputs HTML content for an email instead of pulling it from the approved Postmark template catalog.

✓ INSTEAD

Use `list_email_templates` to find the canonical template ID, and then use this MCP to ensure your agent sends emails using only the official version.

Ignoring server configuration changes.

✗ AVOID

The team updates its email sending structure but forgets to verify if the correct servers are active in the Postmark account settings.

✓ INSTEAD

Run `list_account_servers` and `get_server_info` before any major send campaign. This confirms you're connected to the right endpoint.

The Right Fit

Use this MCP if your primary concern is transactional email delivery reliability, template management, and detailed bounce analysis. You need a single source of truth for all Postmark data—sending history, performance metrics (`get_delivery_stats`), and failure reports (`list_bounces`). Don't use it if you are managing general user relationship communications or internal team chat; those require different connectivity tools (like CRM connectors). If your goal is just to write a simple email draft, writing it locally is fine. But when the

stakes involve compliance, large-scale messaging, or needing verifiable proof of delivery and failure patterns, this MCP is essential.

Checking Email Health Feels Like Juggling Dashboards

Today, checking if a campaign worked means logging into the Postmark UI to see stats, then opening Google Sheets to track bounces, and maybe jumping over to your internal dashboard just to check template IDs. You spend 20 minutes copy-pasting numbers across three different tabs just to answer one question: 'Did it go out?'

With this MCP, the process changes entirely. Instead of logging in, you simply ask your agent for a delivery report. It uses tools like `get_delivery_stats` and `list_bounces`, pulls all that data into one place, and gives you a clean answer instantly.

Sending Emails with Postmark MCP

Manual sending involves selecting the correct template from a dropdown menu, manually inputting recipient lists, and hitting 'send.' If that list is big or if you need to send multiple variations, you're limited by the platform interface.

Now, your agent handles it. You tell it to `send_batch_email` with specific parameters. It manages the volume, verifies signatures, and confirms delivery status automatically—you just get the result.

Postmark with 11 Tools

These tools let your agent perform every function needed for email operations, from sending messages to checking server configuration.

#	TOOL	DESCRIPTION
01	`get_delivery_stats`	Retrieves a report detailing overall email delivery statistics for a specific period.
02	`get_outbound_stats`	Gets delivery metrics specifically related to emails sent out from your account.
03	`get_template`	Fetches the full content and details for a single, specified email template.
04	`list_bounces`	Retrieves a list of recently failed or bounced emails so you can track deliverability issues.
05	`list_domains`	Lists all verified sending domains that your Postmark account is authorized to use.
06	`list_outbound_messages`	Retrieves a list of messages that have been successfully sent from your account.
07	`list_account_servers`	Provides an overview and listing of all active servers associated with your Postmark account.
08	`list_email_templates`	Generates a comprehensive list of every email template available in your catalog.
09	`send_batch_email`	Sends multiple emails to a group of recipients simultaneously using defined templates.
10	`send_email`	Sends a single, immediate transactional email to one or more recipients.
11	`get_server_info`	Retrieves specific configuration details and status information for your connected Postmark server.

See It in Action

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

U Send a transactional email from support@example.com to john@doe.com with subject 'Reset Password'.

I've sent the email. The message ID is '123-abc-456' and it was successfully accepted by Postmark.

U Show me all email bounces from the last 7 days and identify the main failure patterns.

47 bounces in the last 7 days (1.2% bounce rate). Hard bounces: 28 (invalid addresses), Soft bounces: 14 (mailbox full), Spam complaints: 5. Top offending domain: outlook.com with 12 bounces. 8 addresses were auto-suppressed after repeated failures. Your sender reputation score remains healthy at 97/100. Recommendation: clean 28 invalid addresses from your mailing list to maintain deliverability.

U Send a transactional welcome email to new user sarah@meridian.io using the onboarding template.

Welcome email sent successfully via Postmark. Recipient: sarah@meridian.io. Template: "Onboarding Welcome v3" (ID: tpl_2847). Subject: "Welcome to the platform, Sarah!". Message ID: msg_9f3a2b. Delivery confirmed in 0.8 seconds. The template automatically personalized the greeting, included the quickstart guide link, and set the first follow-up email to trigger in 24 hours.

Frequently Asked Questions

01 **How do I find out if my domain is authorized for sending using Postmark MCP?**

You run the `list_domains` tool. This instantly shows you every verified sending domain linked to your account, so you know exactly what domains are safe to use.

02 **Can I list all templates available in my account with Postmark MCP?**

Yes, the `list_email_templates` tool pulls a full catalog of every template ID and name. This is useful for auditing which versions your team uses.

03 **What's the difference between sending one email and using Postmark MCP?**

Sending one email uses `send_email`. If you need to send many emails, like a newsletter blast, use `send_batch_email` for efficiency.

04 **I need to check the history of failed sends using Postmark MCP.**

Use the `list_bounces` tool. It collects and reports on all recent bounces, allowing you to pinpoint invalid addresses or systemic failure patterns.

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 `"postmark-extended": { "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

Postmark 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 Postmark. 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	Postmark MCP
Server ID	`019dd13f-d308-70c6-8c93-0873ed47119f`
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/postmark-extended.