MCP SERVER | NO CODE | CLOUD HOSTED

Email Forwarding Audit MCP

Identify unauthorized data exfiltration and compliance risks in mailboxes.

Email Forwarding Audit MCP lets your AI client scan mailbox configurations to find security holes. It identifies rules pointing to unapproved external addresses, flags inactive configurations, and catches policy violations that could lead to data leaks. You can use it to keep your organization's email traffic within authorized boundaries.

A+ Quality Score 100/100

audit | email | security | compliance | mailbox

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · OCTOBER 2026
vinkius.com/en/ai-agent-connect/email-forwarding-audit

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.

Email Forwarding Audit MCP

4 tools available
Cloud-hosted on Vinkius

Managing mailbox security means knowing exactly where your data is going. This MCP gives your AI client the ability to inspect forwarding rules across your organization to ensure they follow your security policies. Instead of manually checking individual mailboxes, you can ask your agent to find specific risks, like rules sending mail to personal accounts or unverified domains.

It helps you catch shadow IT or malicious forwarding setups that bypass standard controls. You can use it to clean up old, inactive rules that clutter your environment or to verify that every active rule matches your approved list of destinations. It's a direct way to turn your AI client into a security auditor that can pinpoint high-risk policy violations in seconds.

Core Capabilities

01 — Exfiltration Detection

Your agent uses this to find rules sending data to unauthorized external addresses.

02 — Policy Verification

The AI client runs this to check if active rules violate your specific security mandates.

03 — Configuration Cleanup

Your agent identifies disabled or inactive rules that no longer serve a purpose.

04 — Compliance Auditing

The AI client compares current mailbox settings against your organization's approved destination list.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/email-forwarding-audit** — connect your AI agent in three steps.

01 Connect your preferred MCP-compatible client to Vinkius.

02 Ask your AI client to scan for unapproved email destinations.

03 Review the list of compliance conflicts identified by the agent.

04 Use the findings to remediate unauthorized or inactive rules.

Connect your AI client to Vinkius and start auditing mailboxes immediately.

Built For

This MCP is built for security and IT teams who need to monitor mailbox configurations without manual inspection.

Security Analysts

They use the MCP to hunt for unauthorized data movement via email.

Compliance Officers

They use the agent to verify that email forwarding adheres to regulatory requirements.

IT Administrators

They use it to audit and clean up mailbox rules across the organization.

What Changes When You Connect

01 Detects unauthorized email destinations instantly.

02 Identifies high-risk policy violations through automated cross-checks.

03 Locates inactive rules to reduce configuration clutter.

04 Provides a clear view of compliant versus non-compliant forwarding setups.

Real-World Applications

Data Leak Prevention

Identify employees who have set up rules to forward work emails to personal accounts.

Security Audits

Run a quick check to ensure all active forwarding rules point to trusted domains.

Policy Enforcement

Find specific rules that conflict with your company's data handling policies.

Mailbox Hygiene

Locate and review disabled rules that are no longer needed in the environment.

Patterns to Avoid

Manual mailbox hunting

AVOID

You spend hours logging into individual user mailboxes to see if they have active forwarding rules set up.

✓ INSTEAD

Let your agent scan the whole organization at once using `identify_compliance_conflicts` to find violations instantly.

Ignoring stale configurations

AVOID

You leave hundreds of old, disabled forwarding rules in your environment because checking them is too tedious.

✓ INSTEAD

Use `get_inactive_rules` to quickly pull a list of every rule that isn't currently active so you can clean them up.

Assuming all rules are safe

AVOID

You assume that because a rule exists, it must be following company policy.

✓ INSTEAD

Run `get_unapproved_destinations` to catch any rules that are sending data to addresses not on your authorized list.

The Right Fit

Security analysts and IT admins should connect this MCP if they need to monitor mailbox security without manual inspection. If you don't have control over mailbox configurations or don't need to audit forwarding rules, this isn't for you. The decision comes down to whether you need to automate the detection of unauthorized data movement via email.

Email Forwarding Audit 4 Tools

These tools let your agent scan mailbox configurations to find unauthorized destinations, inactive rules, and policy violations.

#	TOOL	DESCRIPTION
01	`get_unapproved_destinations`	This tool finds all email addresses used in forwarding rules that aren't on your authorized list.
02	`get_approved_rules`	Use this to pull a list of all forwarding rules that currently meet your organization's compliance standards.
03	`get_inactive_rules`	This tool identifies any forwarding rules that are currently disabled or inactive.
04	`identify_compliance_conflicts`	This tool checks active rules against your security policies to find direct violations.

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 the forwarding rules that are currently approved.

The approved forwarding rules are: Rule ID 101 (destination: internal@company.com), Rule ID 105 (destination: partner@trusted.com).

U Are there any active rules pointing to unapproved email addresses?

Yes, there is one unapproved destination being used: attacker@external.com.

U Find any high-risk compliance conflicts.

A high-risk conflict was found: Rule ID 202 is active and points to an unverified external domain.

Frequently Asked Questions

01 **What can this MCP do for my security team?**

It allows your agent to audit mailbox forwarding rules to find unapproved destinations and policy violations.

02 **How does it detect unauthorized forwarding?**

It uses the get_unapproved_destinations tool to find rules that point to addresses not on your authorized list.

03 **Can I find rules that violate my company policy?**

Yes, the identify_compliance_conflicts tool specifically looks for active rules that break your security rules.

04 **Does this work with my existing AI client?**

Yes, it works with any MCP-compatible client like Claude, Cursor, or Windsurf.

05 **Can it find old or disabled rules?**

Yes, the get_inactive_rules tool identifies all forwarding rules that are currently disabled.

06 **How do I find rules that violate security policies?**

You can use the `identify_compliance_conflicts` tool to find active rules that point to unverified destinations or have overly broad conditions.

07 **Can I filter the audit by a specific date range?**

Yes, the `get_approved_rules` tool accepts `minCreationDate` and `maxCreationDate` parameters to define your audit window.

08 **How are unapproved destinations identified?**

The `get_unapproved_destinations` tool identifies any destination address used in existing rules that is not present in the organization's authorized destination catalog.

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 `"email-forwarding-audit": { "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

Email Forwarding Audit 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 Email Forwarding Audit. 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	October 2026
MCP Server	Email Forwarding Audit MCP
Server ID	`01a0ef41-d7eb-7367-95c6-153a88585117`
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/en/ai-agent-connect/email-forwarding-audit.

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