MCP SERVER | NO CODE | CLOUD HOSTED

Robots.txt Generator MCP

Control exactly what web crawlers see on your site.

The Robots.txt Generator creates syntactically perfect instructions for web crawlers, telling search engine bots exactly what parts of your site they can and cannot crawl. It lets you define specific rules based on user-agent types, set crawl delays, and properly list your sitemaps so Google and Bing index your content correctly.

A+ Quality Score 100/100

robotstxt | seo | web-crawler | sitemap | automation | web-development

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026

vinkius.com/mcp/robotstxt-generator

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.

Robots.txt Generator MCP

3 tools available
Cloud-hosted on Vinkius

Controlling how web crawlers interact with your site is crucial for SEO, and this MCP handles that job entirely. You use it to write the robots.txt file, which gives search engine bots like Googlebot specific instructions about navigating your domain. Need to block access to temporary directories while still allowing indexing of product images? This tool lets you set those rules precisely. It manages complex directives like 'Allow', 'Disallow', and 'Crawl-delay' for different bot types. If you run into syntax issues, you can check everything first, ensuring every directive follows the standard protocol before publishing. When you connect this MCP through Vinkius, your agent handles all the complexity; you just tell it what rules you need to enforce.

Core Capabilities

01 — Generate robots.txt file

Creates a complete, correctly formatted robots.txt file based on the specific crawling instructions you provide.

02 — Validate rule syntax

Checks your proposed rules and paths to ensure they are syntactically correct before publishing the file to avoid crawler errors.

03 — Get configuration summary

Provides an audit report on your current robots.txt setup, helping you understand all active directives across your site.

One Click on Vinkius — From Prompt to Execution

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

01 First, tell the MCP which bot (user-agent) needs specific rules and what paths it should follow.

02 Second, run the appropriate tool to validate that all syntax is clean or to generate a summary of existing rules.

03 Finally, use the generation function to output the finished robots.txt content for immediate deployment.

The bottom line is you get a guaranteed compliant file that manages bot access without needing manual knowledge of web crawling standards.

Built For

SEO specialists and content managers need this. They spend hours manually checking robots.txt files for errors, worrying about disallowed pages or missed sitemap pointers. This MCP lets them automate the entire process, guaranteeing compliance every time.

SEO Specialist

Uses this to ensure Googlebot can find all critical content while blocking access to private staging areas.

Web Developer

Integrates this tool into build scripts to validate that new code deployments haven't inadvertently blocked crawlers from essential assets.

Digital Marketing Manager

Runs configuration summaries before a site launch to audit and verify all indexing rules across the entire domain structure.

What Changes When You Connect

01 Stop guessing about indexing rules. Use the `get_configuration_summary` tool to audit your existing setup and confirm every bot directive is working as intended.

02 Guarantee compliance with the `generate_robots_txt` function. Just tell it which paths need blocking or allowing, and it spits out a perfect file for deployment.

03 Catch syntax errors before they hurt SEO. The `validate_rule_syntax` tool checks your rules instantly, preventing manual mistakes that waste crawling budget.

04 Manage multiple directives (Allow, Disallow, Crawl-delay) in one place. You don't need to memorize the entire standard protocol; just describe what you want done.

05 Control bot behavior by user-agent. You can write specific rules for Googlebot versus a specialized analytics crawler, optimizing resource usage.

Real-World Applications

A staging site leaks private content

The dev team needs to block all bots from the `/staging/` directory immediately. They ask their agent to generate a robots.txt file that only uses 'Disallow: /staging/' for all user-agents, ensuring the build is live before going public.

Site structure changes often

After migrating product lines, the SEO team needs confirmation that no critical directories were accidentally blocked. They run `get_configuration_summary` to audit all current rules and confirm the paths are open for search engines.

Adding a new content type

A marketing campaign launches thousands of articles, requiring specific crawl delays. The manager uses the MCP to generate the correct file structure, including 'Crawl-delay' rules, and publishes it instantly.

Testing custom bot behavior

The developer needs to make sure a new API endpoint is inaccessible but also wants to check if the path syntax for that block rule is valid. They use `validate_rule_syntax` first, then confirm with generation.

Patterns to Avoid

Manual comma separation

✗ AVOID

Writing a complex list of rules by hand and forgetting the correct colon or semicolon placement.

✓ INSTEAD

Never try to manually write out your instructions. Use the MCP's dedicated tools: run `validate_rule_syntax` first, then let it handle the file creation using `generate_robots_txt`.

Confusing global rules with specific ones

✗ AVOID

Writing a general rule that blocks an entire category of pages when only one subdirectory needs blocking.

✓ INSTEAD

Don't use vague language. Be precise: define the user-agent, and then use the appropriate tool to generate the file content while specifying the exact path for 'Disallow'.

Assuming syntax is correct

✗ AVOID

Publishing a robots.txt file that looks right but contains a minor protocol error (e.g., missing a trailing slash).

✓ INSTEAD

Always run the `validate_rule_syntax` tool before you generate or deploy anything. It catches these small, critical mistakes.

The Right Fit

Use this MCP if your problem is structural: you need to write or audit a text file that controls web crawler access using official standards (Disallow/Allow directives). You should use it when deciding which bots see what content. Don't use it if your goal is to manage internal user permissions, handle database records, or send emails—those require different tools entirely. If you simply need help writing an article about SEO best practices, this isn't the tool; but if you need to tell Googlebot *not* to crawl your test folder at all, this MCP handles it perfectly.

Web crawlers are a nightmare of forgotten rules.

Today, updating robots.txt is a tedious process. You usually have to jump between multiple dev environments and content management systems, manually checking if the changes you made for one bot (like Googlebot) didn't accidentally block another important crawler, or worse, that you forgot to list your sitemap URL entirely.

With this MCP, you simply tell your agent what needs blocking and what needs allowing. The system generates a clean file instantly, guaranteeing all directives are correctly formatted for every bot type—you get perfect control without the manual headache.

Generate robots.txt file: Guaranteed Compliance

You eliminate the need to constantly cross-reference protocol guides or ask a developer to check your syntax for you. You don't have to manually audit whether the 'Allow', 'Disallow', and 'Crawl-delay' rules work together correctly.

It's done. Your agent generates the file, validating every character along the way. This is how you maintain rock-solid SEO compliance.

Robots.txt Generator with 3 Tools

These tools let you generate compliant robots.txt files, check the syntax of your paths, and audit all current web crawler rules in one place.

#	TOOL	DESCRIPTION
01	`generate_robots_txt`	Generates a complete robots.txt file based on agent rules.
02	`get_configuration_summary`	Provides an overview of your current robots.txt setup and active directives.
03	`validate_rule_syntax`	Checks the syntax of any specified robots.txt paths to prevent formatting errors.

See It in Action

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

U **Generate a robots.txt file for all bots that disallows the '/private/' directory and includes a sitemap at https://mysite.com/sitemap.xml.**

User-agent: *
Disallow: /private/

Sitemap: https://mysite.com/sitemap.xml

U **Create rules for Googlebot to allow '/images/' but disallow '/temp/', with a crawl delay of 10 seconds.**

User-agent: Googlebot
Disallow: /temp/
Allow: /images/
Crawl-delay: 10

U **Check if my rules for 'BadBot' are valid: agentName: 'BadBot', disallowedPaths: ['/api/v1/'], allowedPaths: []**

The syntax for the rules provided is valid.

Frequently Asked Questions

01 **How do I use robots.txt Generator to block a folder?**

To block a folder, you must specify 'Disallow:/folder/path/' when generating the file using `generate_robots_txt`. This prevents any bot from crawling that specific directory.

02 **Does robots.txt Generator work for all search engines?**

Yes, it handles rules for multiple user-agents (like Googlebot or Bingbot), ensuring your instructions apply broadly across major web crawlers.

03 **What is the best way to check if my robots.txt file has errors?**

Run the `validate_rule_syntax` tool first. This function checks for protocol and syntax mistakes, giving you confidence before deployment.

04 **Can I use get_configuration_summary with this MCP?**

Yes, running `get_configuration_summary` audits your current settings. It provides a comprehensive overview of all the active directives currently in place on your site.

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 `"robotstxt-generator": { "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

Robots.txt Generator 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 Robots.txt Generator. 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	Robots.txt Generator MCP
Server ID	`019f0488-c376-7144-a77a-c279e5acaeb1`
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/robotstxt-generator.