MCP SERVER | NO CODE | CLOUD HOSTED

Canonical Checker MCP for AI Agents

Fixing Duplicate Content Risks Across All Website Variations

Canonical Checker finds and validates the single, authoritative version of any URL. It detects structural problems caused by inconsistent protocols (HTTP vs HTTPS), differing hostnames (www vs non-www), or variable tracking parameters that muddy your search engine results. Stop worrying about duplicate content risks and ensure all variations point to one clean source.

A+ Quality Score 100/100

seo | canonical | url-audit | duplicate-content | web-standards

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

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.

Canonical Checker MCP

3 tools available
Cloud-hosted on Vinkius

Dealing with complex URLs is a massive headache for SEOs. When multiple versions of the same page exist—like one with 'www' and another without, or one using HTTP while the main site uses HTTPS—search engines get confused. This MCP handles that mess. It analyzes all your URL variations to pinpoint the cleanest canonical version you need. You can use it to check for structural discrepancies across a group of links, separating necessary tracking codes from static page elements. If this is something you deal with often, Vinkius hosts this entire catalog so you connect once and get access to hundreds of specialized tools like this one.

Your agent uses this MCP to manage the critical details of your site structure, ensuring that whether a URL has trailing slashes or different query parameters, you always know which version is the right one for Google.

Core Capabilities

01 — Determine the best canonical tag candidate

The MCP figures out which specific version of a URL should be considered the main source by search engines.

02 — Spot structural discrepancies in groups of URLs

It checks multiple URLs you provide to find differences, like protocol mismatches or slash variations.

03 — Separate static from variable tracking parameters

The MCP analyzes query strings to tell you which parameters are always there (static) and which ones change based on user behavior (variable).

One Click on Vinkius — From Prompt to Execution

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

01 You feed the agent a list of URLs that might be pointing to the same content, regardless of variations.

02 The MCP processes these inputs, checking protocols, hostnames, and query parameters against established web standards.

03 Your agent receives back clear reports identifying structural inconsistencies or providing the single most authoritative URL candidate.

The bottom line is you get a precise audit of your link structure, eliminating guesswork about duplicate content penalties.

Built For

SEO specialists and web developers need this MCP. If your job involves auditing site maps or managing large content inventories, you know how quickly inconsistent URLs can tank rankings. This tool gives you the certainty you need to clean up the messy parts of a website.

SEO Analyst

You use this MCP daily when auditing client websites, running checks to ensure all variations point back to the single correct canonical source.

Web Developer

When implementing site-wide redirects or fixing legacy link structures, you rely on it to guarantee protocol and slash consistency across the entire build.

Digital Marketing Manager

You need it when launching new campaign tracking URLs; this helps separate essential UTM parameters from non-essential session IDs so your data is clean.

What Changes When You Connect

01 Guarantee proper HTTPS usage across your site. Instead of guessing, use the tool to confirm protocol consistency and fix HTTP/HTTPS mismatches.

02 Stop duplicate content penalties before they start. By using identify_canonical_candidate, you ensure search engines know exactly which version is primary.

03 Clean up noisy tracking data. Use detect_parameter_drift to separate critical marketing parameters from random session IDs that clutter your link audits.

04 Streamline site maintenance checks. audit_url_consistency allows you to check a batch of URLs instantly for common structural problems like trailing slashes or www usage.

05 Save time on manual auditing. Instead of checking these variations in multiple tabs, let your agent run the full audit and get one clean report.

Real-World Applications

Website migration from HTTP to HTTPS

A site owner needs to confirm that all old HTTP links are correctly recognized as pointing to the secure, canonical version. The agent runs a full audit and confirms every single link variation now resolves cleanly via HTTPS.

Analyzing multi-platform campaign links

A marketer collects dozens of UTM-tracked URLs from email, social media, and paid ads. Using detect_parameter_drift, the agent immediately flags which parameters are temporary session IDs versus permanent tracking codes.

Standardizing a large e-commerce catalog

A developer must audit thousands of product URLs that vary by trailing slash and hostname (www vs no www). The agent runs a consistency check, providing a definitive list of the standardized URL format.

Fixing internal link rot after a content dump

After uploading hundreds of new articles, the SEO team suspects inconsistent linking. They use identify_canonical_candidate to quickly verify that every article page has correctly identified its single best canonical version.

Patterns to Avoid

Checking URLs manually in a spreadsheet

✗ AVOID

A user copies 50 variations of a URL into Excel and tries to spot all the inconsistencies like 'http' vs 'https', or missing slashes. This is slow, error-prone, and misses hidden parameter issues.

✓ INSTEAD

Feed that entire list directly to the agent. Use audit_url_consistency to check every link against structural rules at once, giving you a clean pass/fail report instantly.

Assuming the canonical version is simple

✗ AVOID

A user assumes the cleanest URL is always the one without 'www', but misses that some industry standards require it. They waste time manually adjusting links based on flawed assumptions.

✓ INSTEAD

Let identify_canonical_candidate do the heavy lifting. It knows current web standards and tells you which specific hostname structure works best for your content.

Ignoring query parameters

✗ AVOID

A developer sees a URL with 'q=shoes' and thinks it's fine, but misses that the site also uses 'sessionid=123' in some instances. They deploy links that are technically valid but confuse search engines.

✓ INSTEAD

Always run detect_parameter_drift first. This separates the essential, static tracking parameters from the variable garbage so you only clean up what matters.

The Right Fit

Use this MCP if your core problem is URL inconsistency: figuring out which version of a page (e.g., www vs non-www, http vs https) should be treated as the main source by search engines. It excels when you have large batches of links to audit or need to differentiate between necessary tracking codes and random session IDs. Don't use it if your problem is something else entirely; for instance, if you just need to rewrite content with AI, that requires a different kind of tool. If you only need basic link checking without worrying about structural standards like query parameters, other simpler auditing tools might suffice, but this MCP covers the full spectrum of web canonicalization issues.

Canonical Checker MCP for AI Agents: Solving Duplicate Content Risks in SEO Auditing

Right now, cleaning up link variations is a painful manual process. You're copying lists of URLs into spreadsheets, manually checking if every link has HTTPS and no trailing slashes. You spend hours trying to figure out if that 'www.' or the 'utm_source=' parameter should actually be there for Google to understand your content correctly.

With this MCP, you give your agent a batch of links. It instantly checks them against best practices, telling you exactly where the protocol fails and which structural elements are causing duplicate signals. You get an actionable report that immediately tells you what needs fixing.

Canonical Checker MCP for AI Agents: Managing Link Structure Complexity in Web Development

Developers often struggle with the sheer volume of link variations, especially when integrating new tracking systems. You have to write complex rules to decide if 'page/item' is the same as 'page/item/' or if a specific session parameter should be stripped out entirely.

Now, you let your agent handle this complexity. It analyzes the structure and gives you clean parameters for development, allowing you to build robust redirects that guarantee every user lands on the intended, single canonical page.

Canonical Checker: 3 Tools for URL Auditing & SEO Consistency

These tools help you audit link structure by checking for canonical candidates, detecting parameter drift, or verifying overall URL consistency.

#	TOOL	DESCRIPTION
01	`identify_canonical_candidate`	Determines the optimal candidate URL to use for a canonical tag.
02	`detect_parameter_drift`	Identifies which query parameters are responsible for causing different versions of a URL.
03	`audit_url_consistency`	Checks and flags structural discrepancies across a provided group of URLs.

See It in Action

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

U I have these three URLs from my campaign: https://mysite.com/product?utm=a&session=123, http://www.mysite.com/product/, and https://mysite.com/product.

Canonical Audit Report

- **Best Candidate:** `https://mysite.com/product` (Uses HTTPS, no 'www', minimal parameters.)
- **Issues Found:** Protocol mismatch (HTTP vs HTTPS) and hostname variation (www). The session ID parameter is variable.
- **Action:** Fix the protocol and standardize the domain to remove 'www'.

U Check consistency for these product pages: https://shop.com/item/, http://shop.com/item, and https://shop.com/item?source=email

Consistency Check Results

- Found 2 structural discrepancies.
- 1. Protocol Mismatch: HTTP vs HTTPS detected.
- 2. Trailing Slash Inconsistency: The first two URLs vary on the trailing slash, but this is minor compared to the protocol issue.

Recommendation: Use the secure HTTPS version consistently.

U Look at these URLs and tell me which parameters are static versus random tracking info: https://mysite.com/search?q=shoes, https://mysite.com/search?q=shirts&utm_medium=social, and https://mysite.com/search?q=hats&sessionid=123

Based on the input links:

- **Static Parameter:** `q` (This parameter reliably identifies the core search term.)
- **Variable Parameters:** `utm_medium`, `sessionid` (These parameters change with every campaign or user visit, and should not be included in canonical tags.)

Frequently Asked Questions

01 **What is the Canonical Checker MCP for AI Agents used for?**

This MCP helps you identify and correct duplicate content risks caused by inconsistent URLs. It ensures all versions of your page—no matter how they are linked—point back to one single, clean canonical source.

02 **How does the Canonical Checker MCP for AI Agents handle HTTP vs HTTPS?**

It automatically checks your entire link group for protocol mismatches. It tells you if some links use insecure HTTP while others use secure HTTPS, helping you standardize to the best practice.

03 **Can I use the Canonical Checker MCP for AI Agents to clean up my tracking URLs?**

Yes. The tool separates static page identifiers from variable tracking codes (like session IDs or campaign tags), so you can keep your marketing data clean while fixing structural errors.

04 **Is the Canonical Checker MCP for AI Agents just for big websites?**

No, it works on any site. Whether you have a small portfolio or an e-commerce giant, this MCP audits protocols, slashes, and parameters to find the single best URL candidate.

05 **What if I use the Canonical Checker MCP for AI Agents and still have issues?**

If you follow its recommendations—like standardizing your hostname or fixing a protocol mismatch—you'll eliminate most of the common duplicate content signals. It gives you the framework to fix 90% of the problem.

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 `"canonical-checker": { "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

Canonical Checker 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 Canonical Checker. 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	Canonical Checker MCP
Server ID	`019f14d9-d8b6-708e-8409-37575aee7744`
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/canonical-checker.