MCP SERVER | NO CODE | CLOUD HOSTED

Paperless-ngx MCP

Your AI agent becomes a digital librarian for your documents.

Paperless-ngx connects your digital archive to any AI agent, turning massive document folders into a searchable knowledge base. Use it to find tax receipts, contract details, or research papers instantly by asking natural questions. Manage tags, upload new files, and retrieve OCR text directly from your existing records.

A+ Quality Score 98.33/100

digital-archive | ocr | full-text-search | file-management | document-indexing | open-source

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

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.

Paperless-ngx MCP

26 tools available
Cloud-hosted on Vinkius

Stop wading through endless folders just to find one date or name. This connector lets you treat your document archive like a librarian that knows every word in it. Instead of manually searching folder structures or running complex queries, you simply ask your agent for what you need—and get an answer based on the full text and metadata of your files. You can use dedicated functions to upload new documents, organize them by creating tags or updating correspondents, or pull out a specific document's details. If you're managing records across different systems, connecting it through Vinkius gives you one place to access this power from any compatible client.

It’s about turning static files into actionable data.

Core Capabilities

01 — Search and Discover Records

Find specific documents by querying the full text, date ranges, or predefined organizational tags.

02 — Manage Document Lifecycle

Upload new files to the archive, download originals for review, or delete outdated records.

03 — Organize Metadata and Tags

Create, update, or retrieve tags and document types to keep your entire library consistently organized.

04 — Inspect Document Content

Fetch the complete OCR text and detailed metadata for any file so your agent can analyze its contents fully.

One Click on Vinkius — From Prompt to Execution

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

01 Subscribe to this MCP and provide the necessary Paperless-ngx API URL and Personal API Token.

02 Connect your preferred AI client (Claude, Cursor, etc.) to Vinkius via the MCP connection point.

03 Directly ask your agent questions about your documents. The agent will call the appropriate tools to search, retrieve data, or execute actions.

The bottom line is you tell the AI what you want in plain English; it handles the complex API calls and retrieval logic for you.

Built For

Anyone dealing with large volumes of paper-based or digital records—legal teams, research academics, or small business operations managers. If your job involves finding specific clauses in old contracts or compiling reports from disparate sources, this is for you.

Legal Paralegal

Retrieving all documents related to a single correspondent over a five-year period and verifying document types.

Operations Manager

Uploading monthly utility bills, creating the correct tag for 'Utilities,' and updating the record's metadata automatically.

Archivist/Records Specialist

Listing all saved views or document types to ensure consistent classification across a historical collection of records.

What Changes When You Connect

01 Find exactly what you need without digging through folders. The `list_documents` tool allows your agent to search across keywords, dates, and tags instantly. No more manual filtering.

02 Consolidate all document data into one place. By using the MCP, your AI client can access every detail of a file via `get_document`, giving you full OCR text for analysis.

03 Maintain perfect order with metadata management. You can use tools like `create_tag` and `update_document` to ensure new files are classified correctly right away.

04 Visualize and manage your library instantly. The agent can provide a quick visual context using `preview_document`, so you don't have to download every file just to check it out.

05 Build robust reporting workflows. By listing all document types (`list_document_types`) or saved views (`list_saved_views`), your AI client builds reports based on consistent, predefined structures.

Real-World Applications

Verifying Contract Dates

A paralegal needs to know the start date of every contract signed by 'Smith & Sons'. They ask their agent to list documents and filter by correspondent. The agent uses `list_documents` and `get_correspondent` to pull up all relevant files, providing a concise timeline instead of a folder dump.

Processing Bulk Invoices

An ops manager receives 50 invoices. Instead of manually uploading them, they use the agent to upload_document in batches. The agent then uses `create_tag` and `update_document` immediately after ingestion, ensuring every file is categorized as 'Invoice' and tagged with the correct vendor name.

Researching Historical Data

A researcher needs all documents related to 'Project Chimera' from 2018. They ask their agent for a list of documents, specifying both keywords and the date range. The agent calls `list_documents`, returning only the relevant records so they can start writing immediately.

Updating Client Records

A legal team updates client information. Instead of logging into two different systems, they tell their agent to update the correspondent details using `update_correspondent`, ensuring all archival files linked to that person reflect the current contact info.

Patterns to Avoid

Treating it like a file share

✗ AVOID

Asking your agent, 'Find me the folder for 2023 tax returns.' The agent can't navigate folders; it only speaks to structured data and metadata.

✓ INSTEAD

Instead, ask: 'List all documents related to taxes from 2023 using `list_documents`. Then use `get_document` on the ID you want details for.' Always reference criteria, not folder paths.

Manual Data Entry

✗ AVOID

Having to manually read a document and type out its key information (e.g., client name, date) into a spreadsheet.

✓ INSTEAD

Use the agent to call `get_document` on the file ID. The full OCR text is returned, allowing your AI client to extract and structure the data points for you automatically.

Ignoring existing categories

✗ AVOID

Treating every document as new and not giving it any context or type.

✓ INSTEAD

Before uploading, use `create_document_type` to define what a 'Warrant' is. Then, when you upload the file using `upload_document`, the system automatically categorizes it correctly.

The Right Fit

Use this MCP if your primary pain point is finding or organizing information locked inside structured documents—contracts, invoices, reports. If you need to know *what* happened based on a document's content (e.g., 'What was the monthly rent mentioned in this lease?'), this is ideal because of tools like `get_document` and OCR text retrieval.

Don't use it if your data lives outside Paperless-ngx, or if you are trying to summarize general knowledge that has nothing to do with a specific file archive. If your goal is simply conversational chat without document grounding, just use a standard LLM. This MCP adds the critical layer of 'document truth' to your conversations.

The Problem: Document searches are always too slow.

Today, finding one piece of information means navigating multiple systems or running complex queries in a dashboard. You have to remember if the record is filed under 'Client' or 'Project,' and then manually cross-reference dates across different folders just to confirm who signed what.

With this MCP, your agent acts like an internal expert. You ask it directly: 'Who was responsible for the billing in March 2023?' The agent doesn't browse; it uses its tools—like `list_documents` and filtering by date—to pull the specific records and deliver the answer instantly.

Paperless-ngx: Full document lifecycle management with Paperless-ngx MCP

The tedious steps of manually creating tags, then having to log back in later to update a correspondent's name on several records are gone. You simply ask the agent to `update_correspondent` or run `create_tag`, and it handles the API calls across your entire archive.

Your AI client now manages the entire document lifecycle for you, from initial upload via `upload_document` to final retrieval using `get_document`. It's a seamless transition from paper chaos to conversational data access.

Paperless-ngx: 26 Tools for Document Management

These tools give your agent the specific actions it needs to manage documents, metadata, correspondents, and tags within your Paperless-ngx archive.

#	TOOL	DESCRIPTION
01	`list_tags`	Provides a complete list of every tag currently in use in the system.
02	`delete_tag`	Removes an unused tag from your system.
03	`download_document`	Retrieves the original physical file for a selected document ID.
04	`get_correspondent`	Fetches all details about an existing correspondent by name or ID.
05	`get_document`	Retrieves the full metadata and OCR text for a specific document file.
06	`get_document_type`	Fetches details about a defined document type category.
07	`get_tag`	Retrieves the full definition of a tag by its ID.
08	`list_correspondents`	Shows a list of all correspondents currently tracked in your archive.
09	`list_document_types`	Displays all the document types and categories you have defined.
10	`list_documents`	Searches and lists documents across the entire archive, supporting filters by date or keyword.
11	`preview_document`	Generates and shows a visual preview link for a selected document.
12	`thumb_document`	Creates and retrieves a small thumbnail image for a document.
13	`update_correspondent`	Modifies the details of an existing correspondent record.
14	`update_document`	Changes the metadata (like tags or date) associated with a specific document.
15	`update_document_type`	Modifies the definition of an existing document category.
16	`update_tag`	Updates the description or name of a tag.
17	`upload_document`	Ingests and processes a new physical file into your document archive.
18	`create_tag`	Generates a new keyword or tag to classify documents.
19	`create_correspondent`	Adds a new individual or company name used in your archive.

#	TOOL	DESCRIPTION
20	create_document_type	Defines and adds a new category for documents, such as 'Invoice' or 'Lease Agreement'.
21	create_saved_view	Saves a specific filter combination so you can easily recall that view later.
22	list_saved_views	Lists all predefined filter views that you have saved for quick access.
23	delete_correspondent	Permanently removes an individual or company from your correspondent list.
24	delete_document	Removes a specific document file from the archive.
25	delete_document_type	Deletes an existing document category definition.
26	delete_saved_view	Removes a previously saved filter view.

See It in Action

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

U Search for all documents related to 'Electricity Bill' from 2023.

I've searched your archive using `list_documents`. I found 3 documents: 'Jan 2023 Electricity' (ID: 101), 'Feb 2023 Electricity' (ID: 105), and 'March 2023 Electricity' (ID: 110). Would you like to see the details of one of them?

U Upload a new document titled 'Contract 2024' with tag ID 12.

Processing the upload... I've successfully used `upload_document` to add 'Contract 2024'. It has been assigned ID 250 and is now being processed by Paperless-ngx for OCR.

U Get the full content and a preview of document ID 42.

I've retrieved the data for document 42 using `get_document` and `preview_document`. The document is a 'Lease Agreement'. The OCR text mentions a monthly rent of $1,200. I've also generated a preview link for you.

Frequently Asked Questions

01 **How do I search my documents with Paperless-ngx MCP?**

You use the agent and ask it to list documents. The underlying tool, `list_documents`, allows you to specify filters like dates or keywords in your prompt.

02 **Can I upload new files using the Paperless-ngx MCP?**

Yes. You use the `upload_document` tool. The agent handles sending the file and ensuring it gets processed by Paperless-ngx for OCR indexing.

03 **What if a document needs more tags? Can I update them with Paperless-ngx MCP?**

You can use `update_document` to change metadata, or you can call the specific `create_tag` tool if the tag doesn't exist yet.

04 **Do I need API tokens for Paperless-ngx MCP?**

Yes. You must provide your Personal API Token and API URL when subscribing to this MCP so the agent can authenticate and talk to your instance.

05 **Can I retrieve just the text from a document using Paperless-ngx MCP?**

Absolutely. The `get_document` tool retrieves complete OCR text, allowing your AI client to read and summarize the content without needing a visual preview.

06 **Can I search for documents using specific tags or date ranges?**

Yes. The `list_documents` tool allows you to filter by `tags__id__in` and `created__date__gte`. You can also perform a full-text search using the `query` parameter.

07 **Is it possible to see a preview of a document without downloading the whole file?**

Absolutely. Use the `preview_document` or `thumb_document` tools to get visual representations of the document content directly through the agent.

08 **Can I create new organization categories like tags or correspondents via AI?**

Yes, you have full management capabilities. You can use `create_tag`, `create_correspondent`, and `create_document_type` to organize your archive on the fly.

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 `"paperless-ngx": { "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

Paperless-ngx 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 Paperless-ngx. 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	Paperless-ngx MCP
Server ID	`019e38d2-4914-73e8-a4af-17ec62c81c35`
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/paperless-ngx.