MCP SERVER | NO CODE | CLOUD HOSTED

Cradl AI MCP for AI Agents

Extract Key Data Points from Invoices, Receipts, and Custom Forms

Cradl AI equips your agent to read and structure data from any document type, whether it's a complex invoice, a simple ID scan, or a custom form. It uses deep learning models to pull out key details—like dates, amounts, names, and IDs—and turn them into clean, usable text for your workflow.

A+ Quality Score 100/100

ocr | data-extraction | deep-learning | document-processing | automation | structured-data

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

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.

Cradl AI MCP

10 tools available
Cloud-hosted on Vinkius

Dealing with documents is messy work. You get PDFs, scans, JPEGs, and forms that look different every time, making manual data entry a nightmare. Cradl AI changes that by letting your agent send the document URL directly to the system. It uses deep learning models to analyze the file and pull out exactly what you need—like invoice numbers or customer names—into structured, actionable fields.

It's built specifically for high-volume data processing in finance and operations. You don't write complex parsers; you just point your agent at the document, and it does the heavy lifting. If you already use Vinkius to connect various services, adding Cradl AI keeps all your document intelligence centralized right where your agent works.

After extraction, you can track everything—from listing available models to checking if a large batch of documents finished processing. It turns unstructured paper trails into clean data points that your application can use immediately.

Core Capabilities

01 — Extract Data from URLs

Pulls structured key-value pairs directly from the content of any document hosted online.

02 — Check Task Status and Results

Confirms if a specific document processing job is finished, and provides extracted fields along with confidence scores.

03 — Manage Document Processing Batches

Lists all processed document batches or retrieves detailed summaries for an entire group of files.

04 — List Available Extraction Models

Shows the names, versions, and training status of every custom-trained data extraction model you own.

05 — Review Workflow Settings

Retrieves the structure and configuration for specific document processing flows.

One Click on Vinkius — From Prompt to Execution

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

01 First, your agent sends a document URL to this MCP. The system runs an OCR engine that processes the image or text.

02 Next, the deep learning model attempts to predict and normalize data boundaries based on the file type (e.g., recognizing an invoice number vs. a total amount).

03 Finally, your agent receives clean JSON output containing all the extracted key-value pairs, ready for immediate use.

The bottom line is: you send it a document link, and you get structured data back without manual cleanup or coding boilerplate.

Built For

If your job involves reading documents that don't fit neatly into spreadsheets—think invoices, receipts, ID scans, or old contracts—you need this. It targets the operations manager tired of manual data entry and the finance analyst who needs rapid visibility into large batches of financial records.

Financial Analyst

Uses the MCP to process hundreds of receipts or invoices daily, instantly populating ledger entries without human intervention.

Operations Manager

Monitors high-volume document submission queues, checking task statuses and batch results to ensure processing accuracy across departments.

Developer

Tests model performance by listing available models or auditing workflow settings to integrate robust data pipelines into client applications.

What Changes When You Connect

01 Instead of manually reading PDFs, you send the document link once. The system handles the deep learning analysis and returns clean data fields instantly.

02 You can monitor high-volume operations using tools like `list_batches` or `get_task_status`, giving immediate visibility into processing success rates.

03 Developers gain confidence by running checks on model performance via `get_model_details` before deploying the MCP to production workflows.

04 The ability to list available extraction models means you always know what parsing capabilities your agent has access to, minimizing integration guesswork.

05 When a process fails, tools like `list_processing_tasks` quickly locate the failed task ID and status, cutting down on debugging time.

Real-World Applications

Processing Quarterly Expense Reports

An operations manager receives 50 vendor expense reports. Instead of opening each PDF to copy dates and amounts, the agent uses the MCP to send all URLs in a batch. The system extracts every required field, giving the manager one clean spreadsheet ready for accounting.

Onboarding New Employees

A HR specialist needs to process multiple employee IDs and contracts. They feed the document links into the agent. Using specialized models, the system extracts names, dates of birth, and ID numbers accurately, allowing the onboarding workflow to continue without delays.

Reconciling Supplier Invoices

A finance analyst receives a mix of invoice formats from different suppliers. The agent uses Cradl AI's extraction tools to read every unique document type, normalizing the data structure so it can be automatically uploaded into the accounting platform.

Auditing Historical Documents

A compliance officer needs to verify a year's worth of operational documents. They use the MCP to list workflows and check batch details, ensuring every document in the archive passed through the correct processing steps.

Patterns to Avoid

Assuming text is always clean

✗ AVOID

Trying to copy data from a scanned PDF or an image-based receipt directly into a spreadsheet and hoping it works. This often results in jumbled characters or missing key fields.

✓ INSTEAD

Always use the `extract_data_from_url` tool. This forces the system to run OCR and deep learning prediction, ensuring that even poor quality scans are converted into reliable text data.

Ignoring model performance

✗ AVOID

Using a general-purpose extraction tool without knowing if it's trained specifically on your company's unique invoice format. This results in low confidence scores and missing fields.

✓ INSTEAD

Before running a large job, use `list_extraction_models` and then `get_model_details`. Confirm the model you select was custom-trained for your specific document type.

Losing track of multiple jobs

✗ AVOID

Sending 20 documents to be processed over several hours and forgetting which job failed or when it succeeded. This requires manual checking across various dashboards.

✓ INSTEAD

Use `list_batches` first to get the batch ID, then use `get_task_status` with that specific task ID. This gives a single source of truth for all your document processing results.

The Right Fit

You should connect Cradl AI if your primary pain point is turning unstructured documents into structured data. Use this MCP when you need to read invoices, IDs, or receipts from diverse formats and require high accuracy. Don't use it if all your source files are already clean CSVs; for that, a simple file upload tool works better. However, if you need to manage the complexity of different document types (e.g., separating invoice data from tax ID data), then list models or check flows first. If you only ever deal with one perfect PDF format, maybe a simpler single-purpose connector is fine. But for real-world operations involving multiple sources and formats, Cradl AI provides the necessary control through tools like `get_flow_details` to ensure consistency.

Cradl AI MCP: Automating Invoice and Receipt Data Extraction

Today, finance teams spend hours handling invoices. They download a batch of PDFs or scans; then they open them one by one, manually copying the vendor name, invoice number, total amount, and payment terms into accounting software. This process is slow, prone to typos, and frankly, it's miserable.

With this MCP, you simply point your agent at the folder containing the documents. The system handles the deep learning model execution in the background. You get a single output: structured JSON data containing every needed field from all those PDFs—ready for immediate import.

Cradl AI MCP: Tracking High-Volume Document Processing Status

When running large, multi-day document processing jobs, tracking progress is a nightmare. You have to jump between task IDs and batch numbers just to see if the job succeeded or failed and why.

This MCP centralizes that visibility. You can use `list_batches` to see what's been processed overall, then drill down with `get_task_status` to know exactly where any specific document stands in the workflow—completed, pending, or failed.

Cradl AI: 10 Tools for Structured Document Data Extraction

Use these tools to manage document batches, check task statuses, list available extraction models, or pull data directly from URLs.

#	TOOL	DESCRIPTION
01	`extract_data_from_url`	Triggers a new prediction, using the OCR engine and deep learning models to pull structured data from any provided document URL.
02	`get_batch_details`	Retrieves detailed summaries and status reports for an entire group of processed documents.
03	`get_flow_details`	Shows the specific structure and configuration settings for a designated document processing workflow.
04	`get_model_details`	Retrieves detailed metadata, accuracy metrics, and schema definitions for any trained extraction model.
05	`get_task_status`	Checks the current status of a document task and resolves confidence scores alongside extracted key-value pairs.
06	`list_batches`	Provides a list of all processed document batches, including creation dates and total counts.
07	`list_workflows`	Lists every defined document processing flow, showing their associated triggers and steps.
08	`list_extraction_models`	Retrieves a comprehensive list of all data extraction models available in the system, noting their status and version.
09	`list_processing_tasks`	Lists recent document processing tasks, showing IDs, statuses (e.g., FAILED), and timestamps.
10	`search_models_by_name`	Searches for specific extraction models using a keyword query against the model metadata.

See It in Action

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

U I just got a pile of vendor invoices. Can you extract all the invoice numbers and total amounts for me?

✅ **Extraction Complete.**

We found 12 documents, and here is the structured summary:

INVOICE ID	VENDOR NAME	TOTAL AMOUNT	DATE FOUND
INV-9032A	Acme Corp	$452.11	2024-06-12
REC-77B	Beta Services	$89.00	2024-06-12
INV-1221X	Global Supply	$1,230.55	2024-06-11

U The processing job for the HR documents seems stuck. What's its current status?

Status Check: 🟡 **PROCESSING**

The task is currently running and has processed 85% of the required fields. The system predicts a completion time in about 5 minutes.

- **Total Documents:** 20
- **Completed:** 16 (98% confidence)
- **Failed:** 4 (Check workflow settings for root cause)

Let me know if you want me to check the flow details next.

U I need to find out what kinds of custom models are available for my account.

🔎 **Models Found:**

You currently have access to 4 specialized extraction models:

- **Invoice Parser v3.1:** Best for multi-page billing statements.
- **ID Document Scanner:** Optimized for US/EU government IDs (high accuracy).
- **Receipt Extractor v2.0:** Ideal for quick, low-detail vendor receipts.
- **Custom Form Auditor:** Trained on your specific internal forms.

Which one should I use next?

Frequently Asked Questions

01 **How does Cradl AI help me read scanned receipts and invoices?**

It uses deep learning models to process images, not just digital text. You send the document link, and it performs OCR (Optical Character Recognition) to pull out key details like amounts, dates, and vendor names from poor-quality scans or photos.

02 **Can I use Cradl AI for multiple types of documents? For example, IDs and invoices?**

Yes. You can train and list multiple specialized models (like an ID scanner model separate from an invoice model). This lets your agent select the right tool for every different document type you encounter.

03 **If I process a large batch of documents, how do I know which ones failed?**

You can use the MCP to list all processed batches and then check individual task statuses. This gives you precise feedback on exactly which document caused an error and why.

04 **Is Cradl AI just for finance, or can it handle other types of forms?**

It's not limited to finance. While invoices are a core strength, the system is designed for custom data extraction. You can train models on virtually any structured document—HR onboarding forms, legal contracts, etc.

05 **What should I do if my current model isn't working well on new documents?**

Check the model details using the MCP to review its accuracy metrics. If performance is dipping, you can use the platform's features to audit and improve that specific extraction model.

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 `"cradl-ai": { "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

Cradl AI 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 Cradl AI. 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	Cradl AI MCP
Server ID	`019d757d-923b-710f-8e88-4982a57ea554`
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/cradl-ai.