

MCP SERVER

NO CODE

CLOUD HOSTED

Design Pickle MCP for AI Agents

Manage Brand Guidelines and Production Queue Status

Design Pickle manages your entire creative production process through your AI agent. Connect this MCP to track brand profiles, submit new design requests, and monitor the status of every asset in your queue—all from a single conversation.

A+ Quality Score 100/100

graphic-design

creative-workflow

brand-management

production-queue

design-requests

outsourcing



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.

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.

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.

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

Encryption at rest

Ed25519

PKI vault signatures

24h TTL

Ephemeral session keys

V8 Isolate

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.

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.

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.

06 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot 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.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

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.

Design Pickle MCP

10 tools available

Cloud-hosted on Vinkius

Forget logging into multiple dashboards just to check on an image or confirm a color code. This connector brings Design Pickle's unlimited graphic design service directly into your agent's workflow. You can now ask your AI client questions like, "What's the status of the new product banner?" and get an instant answer with details about the estimated delivery time.

You manage everything from brand guidelines to finished assets using natural language. Need to kick off a new logo concept? Your agent handles creating that request, ensuring it has all the right parameters before sending it over. You can check which logos are available or audit the primary color palette against your existing designs. When you're ready to track history, your agent pulls up every past request and its final delivery date. Because this MCP is hosted on Vinkius, you get access to Design Pickle's full library of tools, letting your AI client handle complex tasks like listing all available request types or even canceling an outdated design job. It's about keeping your entire creative pipeline conversational.

Core Capabilities

01 — Manage the Request Lifecycle

Create new designs with specific instructions and track their progress from submission to completion.

03 — Monitor Production Status

Retrieve a comprehensive list of all design requests, showing their current status, who is assigned to them, and expected delivery timelines.

05 — Update or Finalize Designs

Change the status of a design job, marking it as complete or updating details mid-process.

02 — Audit Brand Guidelines

Access detailed brand profiles, including mandated logos, color palettes, and typography rules for consistency checking.

04 — Search and Retrieve Assets

Quickly find specific past or active designs by title or retrieve detailed information for any given request ID.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/mcp/design-pickle — connect your AI agent in three steps.

- 01** Connect your AI client to this MCP and authorize access using your Design Pickle API Key.
- 02** Tell your agent what you need—for example, 'List all design requests from Q3' or 'Show me the brand color palette.'
- 03** Your agent calls the appropriate tool, pulls the data (like request details or brand guidelines), and presents the structured information back to you in plain conversation.

The bottom line is that your AI client talks directly to Design Pickle's system, giving you real-time status updates without opening a browser tab.

Built For

Creative Directors and Marketing Managers use this MCP when they need constant visibility into their design production queue. They struggle with manual follow-ups and ensuring brand guidelines are followed across multiple campaigns. This tool lets them manage the entire creative lifecycle from a single chat interface.

Marketing Manager

Checks the current status of banner ads or social media graphics for an upcoming campaign, and initiates new requests when needed.

Creative Director

Audits a client's brand profiles against design standards to ensure logos and color palettes are used correctly across all departments.

Small Business Owner

Keeps tabs on the backlog of printed or digital assets, accessing delivered files without logging into a complex portal.

What Changes When You Connect

- 01** Eliminate dashboard fatigue. Instead of logging into a separate portal to check status, you ask your agent about the 'Q3 Campaign Logo' and get immediate details.

-
- 02** Maintain brand consistency instantly. By accessing brand profiles via `list_brands`, you ensure every new design request adheres strictly to approved color codes and logos.
-
- 03** Speed up content creation. Use `create_design_request` to submit a full set of instructions—including required formats and types—without leaving your chat window.
-
- 04** Never lose track of assets again. You can search for old designs using `search_design_requests`, pulling up the specific files you need instantly.
-
- 05** Know exactly where every job stands. Use `list_design_requests` to get a real-time overview of everything in progress, who's assigned, and when it will ship.
-

Real-World Applications

A designer needs to check the status of 15 assets.

Instead of opening 15 different tabs, the agent runs `list_design_requests` for the 'Holiday Campaign' and instantly provides a list of all items, noting which are in progress versus delivered.

A director needs to verify brand compliance.

The user asks, 'What's our official primary color?' The agent calls `list_brands` and provides the exact hex codes and typography rules needed for the next project.

A marketing manager needs to submit urgent creative assets.

The user tells their agent to create a new banner set. The agent uses `create_design_request`, automatically filling out the request type and ensuring all necessary directions are passed in JSON format.

An asset is wrong and needs changing.

If a design request has outdated text, the user doesn't need to email anyone. The agent uses `update_design_request` with the new data, modifying the job directly in the system.

Patterns to Avoid

Treating it like a simple file search

X AVOID

Just asking, 'Where is my banner?' This only gets you filenames and doesn't tell you if the designer has even started working on it.

✓ INSTEAD

Ask your agent to ``search_design_requests`` by keyword, then follow up with ``get_design_request_details``. This gives you the status *and* the file location.

Manually checking guidelines

X AVOID

Constantly switching between a style guide PDF and your project to verify if the right logo version was used.

✓ INSTEAD

Ask the agent to ``list_brands``. It pulls up all mandated logos, color palettes, and typography rules directly in the chat for immediate reference.

Overlooking job completion

X AVOID

Assuming a design is done because the designer said 'it's ready.' You need confirmation it was actually marked as finished.

✓ INSTEAD

After reviewing all assets, use ``complete_design_request`` to formally mark the entire batch as finished in the system record.

The Right Fit

Use this MCP if your primary pain point is visibility across a high volume of creative work. If you need to submit new jobs, check brand guidelines, or track status updates without logging into multiple web portals, this tool is essential. Don't use it if you only need general marketing advice; that requires a content generation model. Also, don't rely on it for final graphic design—it manages the *process* of getting the design done. If your workflow involves highly complex, non-standardized approvals (like legal sign-offs outside of Design Pickle), you'll need an additional custom integration layer to handle those specific steps.

Design Pickle MCP: Managing Creative Production Requests in Marketing

Today, managing a single campaign's assets involves jumping through hoops. You check the status of your logo via one dashboard, confirm the color codes against another style guide, and then submit the job details into a third system. This manual process is slow, prone to copy-paste errors, and makes it impossible for anyone to get a full picture of where everything stands.

With this MCP, you simply tell your agent what you need—'I need 5 social media assets.' Your agent handles the entire sequence: listing available request types, ensuring brand guidelines are checked against the input, and submitting the job using `create_design_request`. You get a confirmed ticket number and an immediate status update.

Design Pickle MCP: Auditing Brand Profiles and Guidelines for Creative Directors

Before sending any asset to design, creative directors must manually pull up the brand guidelines document. Then, they have to cross-reference that PDF against what's currently in use across different departments—is it the right version of the logo? Is the typeface correct?

This MCP lets you audit everything instantly. You ask your agent about 'Main Brand Guidelines,' and it pulls up all required assets and rules using `list_brands`. It gives you the single source of truth, saving hours of cross-departmental email chasing.

10 Tools for Managing Brand Assets and Design Requests

Use these tools to submit new designs, check status updates, audit brand guidelines, or retrieve details about every asset in your account.

#	TOOL	DESCRIPTION
01	<code>complete_design_request</code>	Marks an existing design job as finished within the system.
02	<code>create_design_request</code>	Submits a brand new design request using specific instructions and preferred file types.
03	<code>get_design_request_details</code>	Pulls all detailed information, including status and assignment, for one specific design request.
04	<code>get_request_type_details</code>	Retrieves specific details about a predefined category of design work (e.g., 'Social Media Ad').
05	<code>list_brands</code>	Lists all brand profiles configured for your account, including necessary guidelines.
06	<code>list_request_types</code>	Retrieves a list of all available categories or types of design work you can request.
07	<code>list_design_requests</code>	Generates a master list of every graphic design job in your account, showing status and assignment details.
08	<code>search_design_requests</code>	Finds specific design requests by using keywords located in their title or description.
09	<code>update_design_request</code>	Modifies an existing design request with new instructions or status updates.
10	<code>cancel_design_request</code>	Cancels a specific design job if it is still in a cancellable state.

See It in Action

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

U Can you list every design request I have running for Q3?



Q3 Campaign Requests

- **Social Media Banner Set:** In Progress (Assigned: John D.) - Est. Delivery: Oct 15th
- **Product Info Sheet:** Ready for Review (Status: Needs client sign-off)
- **New Logo Concept:** Pending Triage (ETA: Within 24 hours)

You can ask me to check the specific details or update any of these jobs.

U I need a new banner for our fall sale, but I don't know what type it is.



Sure thing. Based on your goal, you probably want 'E-commerce Promotional Banner.' This is one of the available types. Here are some details:

- **Type ID:** 405
- **Required Input:** Product Line Name, Dimensions (e.g., 1200×628)
- **Example Use:** Perfect for seasonal sales or product launches.

Should I go ahead and create a draft request using this type?

U What's the status of 'Website Hero Image'? And what are our brand colors again?



Okay, I have two answers for you:

1. Website Hero Image Status: It was marked as complete last week and the final file is available in your asset library.

2. Primary Brand Colors: Your core palette includes: #3498db (Primary Blue), #2ecc71 (Success Green), and #f1c40f (Accent Yellow). These are the colors you should stick to.

Frequently Asked Questions

01 How can I use Design Pickle MCP to track my design requests?

You can get a full overview of all your graphic designs by asking your agent to list them. It provides the current status, who is working on it, and when you can expect the final files.

02 What if I need to submit a brand new design job?

You simply tell your agent what you want—like 'I need 5 banners for fall.' The agent will guide you through submitting the request, making sure all required details are captured before it goes live.

03 Can I check brand guidelines using Design Pickle MCP?

Yes. You can ask your agent to list your configured brands. It pulls up the official logos, color codes, and typography rules so you always know what's correct for consistency.

04 Is Design Pickle MCP better than using a simple spreadsheet?

A spreadsheet is static; this MCP gives you real-time status updates. It connects directly to the design service, meaning it knows if the job was marked 'complete' in the system, not just what you typed into a cell.

05 How do I find old designs or completed assets?

You can ask your agent to search through all past requests using keywords from titles. This retrieves detailed information and helps you locate specific finished files quickly.

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 `"design-pickle": { "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

Design Pickle 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](https://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 Design Pickle. 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	Design Pickle MCP
Server ID	019d7584-5ad8-705a-945f-08e3e35c3ed5
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/design-pickle.