MCP SERVER | NO CODE | CLOUD HOSTED

Midjourney MCP

Generate and refine professional AI art in conversation.

Midjourney AI (Generative Image Arts) lets you generate professional, high-fidelity artwork directly from your agent. Use text prompts for new scenes, or feed it existing images to reverse engineer concepts and create variations. It handles complex visual tasks like blending multiple sources, upscaling specific parts of a grid, and simulating camera movements, giving designers full control over their artistic output.

A+ Quality Score 100/100

generative-art | text-to-image | ai-design | creative-tools | upscaling

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/midjourney-ai-generative-image-arts

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.

Midjourney AI (Generative Image Arts) MCP

10 tools available
Cloud-hosted on Vinkius

This MCP gives you complete command over Midjourney's generative art tools, letting your agent handle the heavy lifting of professional image creation. Instead of wrestling with complex web interfaces or endless manual prompting, you talk to it conversationally and get results. You can generate entire scenes from a simple text description, but you don't stop there. If you love a particular angle in an initial 2x2 grid, you can pull out just that tile for high-resolution upscaling. Need to iterate on a style? Just tell your agent to make variations or blend two completely different source images into one cohesive piece. The whole process, from generating the concept to refining the final asset, runs through your preferred AI client and is managed by Vinkius, making advanced creative tools feel simple. You just describe what you need, and it manages the job status until the image is ready.

Core Capabilities

01 — Generate images from text

You provide a description, and the agent triggers Midjourney to create a brand new piece of art.

02 — Refine existing visual concepts

The MCP creates structural variations or entirely new versions based on an image you already generated.

03 — Blend multiple source images

You give it a handful of pictures, and the agent merges them into one cohesive, unique composition.

04 — Enhance resolution or zoom

The MCP takes specific tiles from an image grid and professionally upscales them, or simulates camera movements like panning or zooming out a scene.

05 — Determine source prompts

You provide a public image URL, and the agent analyzes it to return four suggested text descriptions of what's in the picture.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/midjourney-ai-generative-image-arts** — connect your AI agent in three steps.

01 Subscribe to this MCP on Vinkius and input your Midjourney API key.

02 Tell your AI client which creative task you need, such as generating a scene or blending images. The agent executes the command and retrieves an initial Job ID.

03 Your agent monitors the job status automatically. When finished, it hands you the final high-resolution image or data package.

The bottom line is that your AI client manages every step of the complex creative process so you never have to manually check job statuses or piece together disparate images.

Built For

This MCP is built for people who live in the visual space. It's for the art director whose budget limits prompt engineers, or the marketer who needs dozens of high-quality assets this week. If your job requires going from a vague idea to multiple finished mockups quickly, you need this.

Digital Designer

They use it daily to prototype visual concepts and iterate on aesthetic directions without manually entering prompts into the Midjourney website.

Creative Director

They command their agent to generate high-quality moodboards or story arc frames, using both blending tools and variation commands to maintain a specific artistic vision.

Marketing Content Creator

They produce photorealistic assets and cinematic backgrounds directly from their terminal, drastically speeding up the process of filling out campaign visual needs.

What Changes When You Connect

01 You bypass manual prompt entry. Simply tell your agent what you want, and it handles the entire generation process using `generate_image`.

02 Need to improve a single detail? Instead of accepting the initial render, use `upscale_image` to isolate and enhance only the tile you like best.

03 Concepting is easier than ever. If you see an image online but don't know how it was made, run `describe_image` on its URL to pull out potential source prompts.

04 Don't settle for one look. Use `generate_variation` or the `blend_images` tool to quickly iterate and combine styles until your vision is perfect.

05 Manage complexity with job tracking. The MCP handles status updates automatically; you just ask it to check progress using `get_job`.

Real-World Applications

Creating a cinematic moodboard

A director needs a set of background images for a sci-fi film. They prompt their agent with 'futuristic city at sunset' (`generate_image`). After getting the 2×2 grid, they use `pan_image` to simulate a sweeping view and then call `blend_images` using three separate source photos to ground the art in reality.

Improving product mockups

A marketing team generates an initial image of a shoe (`generate_image`). They realize the detail needs work, so they use `upscale_image` on the sole tile. Then, if that doesn't look right, they run `reroll_job` to generate another set of options.

Developing visual storyboards

A game designer needs a sequence of shots. They start with an image and use the agent to call `zoom_out_image`, then follow up by calling `generate_variation` on that new, wider frame, ensuring consistency across multiple 'shots'.

Reverse-engineering artistic styles

A designer finds a beautiful piece of concept art online. They feed the URL to the agent and call `describe_image`, which gives them four candidate prompts they can then use in `generate_variation` to replicate the style.

Patterns to Avoid

Trying to generate without tracking

✗ AVOID

The user generates a massive image set and then forgets the Job ID, leading to confusion when trying to upscale or check status.

✓ INSTEAD

Always retrieve the Job ID immediately after calling `generate_image`. If you need to check later, use that specific ID with the `get_job` tool.

Over-blending too many images

✗ AVOID

The user tries to blend six different source photos into one composition, violating the system limit and receiving an error.

✓ INSTEAD

Remember that `blend_images` accepts a maximum of five sources. If you have more than five concepts, try grouping them into smaller sets first.

Assuming variation means better

✗ AVOID

The user runs `generate_variation` multiple times without analyzing the results, wasting credits on poor iterations.

✓ INSTEAD

Treat variations as drafts. Review the initial grid and only run `generate_variation` when you are confident in the core concept; don't just hit the button repeatedly.

The Right Fit

Use this MCP if your goal is generating, refining, or modifying complex visual assets that need multiple passes—like cinematic storyboards, high-end mockups, or artistic moodboards. You should use it when you are dealing with a multi-stage creative process (e.g., generate -> upscale -> vary). Don't use this if all you need is simple text generation or basic image editing; those tasks are better suited for dedicated photo editors. Also, don't rely on the MCP to write your initial prompt—you still need to be a strong conceptual artist first. However, if you struggle with execution speed or remembering the precise sequence of steps (like panning followed by upscaling), this MCP is essential because it manages that entire pipeline conversationally.

Visual asset creation used to feel like juggling five different software programs.

Today, taking a concept from your head to a finished image involves jumping between text editors for prompting, the Midjourney website for generation, and then separate tools just to crop or upscale the final output. You're copying links, checking job statuses on different pages, and manually stitching tiles together.

With this MCP, you talk to your agent once. The tool handles the whole sequence: it generates the concept, tracks the status, pulls out only the necessary tile using `upscale_image`, and lets you refine that asset right from your chat window. You get the final high-resolution render without ever leaving your workspace.

Midjourney AI (Generative Image Arts) gives you total control over every part of the image.

Manual workflows force you to treat images as static endpoints. If you want a sense of motion, you're limited. You can't easily simulate camera movements or combine disparate visual styles

This MCP lets your agent execute advanced directional commands like `pan_image` and `zoom_out_image`. Your creative output gains cinematic depth that was previously locked behind manual, multi-step post-production work.

without professional graphic design software and hours of effort.

Midjourney AI (Generative Image Arts) MCP - 10 Tools

Use these tools to create, enhance, modify, and analyze images using Midjourney's powerful generative models.

#	TOOL	DESCRIPTION
01	`generate_image`	Creates a new image based on a text prompt using Midjourney's generative models.
02	`get_job`	Checks the current status of an ongoing creative job to see if it's pending or complete.
03	`list_jobs`	Retrieves a history of all your recently executed Midjourney prompts and jobs.
04	`describe_image`	Analyzes an image URL and returns four potential text descriptions that might have created it.
05	`blend_images`	Combines two to five specific images into a single, unique artistic composition.
06	`upscale_image`	Selects and increases the resolution of one specific tile from an image grid.
07	`generate_variation`	Creates new, structural artistic options based on a single existing image in a grid.
08	`reroll_job`	Re-runs an identical prompt exactly as before to generate a completely new set of outputs.
09	`pan_image`	Extrapolates the image borders by simulating a camera movement in a specific direction (up, down, left, or right).
10	`zoom_out_image`	Widens the visible perspective of an existing scene by zooming out.

See It in Action

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

U Generate a futuristic cyberpunk city at sunset with neon reflections in 16:9

Generating image... I've triggered the Midjourney imagine task with your prompt and a 16:9 aspect ratio. Your Job ID is 'mj-98765'. I'm monitoring the process. High-fidelity textures and lighting are being computed. I'll provide the 2×2 grid URL as soon as it's ready.

U Upscale the first image from grid ID 'mj-98765'

Upscaling tile 1... I've triggered the high-resolution render for your first choice. This will isolate the image and apply advanced AI detail enhancement. I'll provide the final high-res URL once the upscale job completes.

U Describe the contents of this image: [url]

Analyzing visual concepts... I've retrieved 4 candidate prompts from Midjourney: 1) 'Minimalist architectural sketch in charcoal', 2) 'Brutalist concrete building perspective', 3) 'Abstract geometric shadows on stone', and 4) 'Architectural blueprint style'. Would you like to use any of these to generate a new variation?

Frequently Asked Questions

01 **How do I check if my Midjourney job is done using the Midjourney AI (Generative Image Arts) MCP?**

You use `get_job` to check the status of any running task. Your agent monitors the process, so you just need to ask it for an update and provide the current Job ID.

02 **Can Midjourney AI (Generative Image Arts) create art from a drawing I find online?**

Yes. You send the image URL to `describe_image`, and the tool will return four candidate text descriptions, helping you understand what prompt was used.

03 **What is the difference between generate_variation and reroll_job in Midjourney AI (Generative Image Arts)?**

`generate_variation` creates new structural options based on a specific image within a grid. `reroll_job`, however, re-runs the exact same prompt arguments to give you an entirely fresh set of results.

04 **Does this MCP support blending photos from different sources?**

Yes, you can use `blend_images` to merge two to five explicit source images into a single unique composition, bridging different artistic styles flawlessly.

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 `"midjourney-ai-generative-image-arts": { "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

Midjourney AI (Generative Image Arts) 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 Midjourney AI (Generative Image Arts). 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	Midjourney AI (Generative Image Arts) MCP
Server ID	`019d75d4-9f1e-7225-a421-3cbfed42b9ed`
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/midjourney-ai-generative-image-arts.