MCP SERVER | NO CODE | CLOUD HOSTED

Stability AI MCP

Generate, refine, and edit professional visuals instantly.

Stability AI gives your agent access to professional image generation tools. You can create photorealistic images from text prompts, upscale low-resolution graphics while keeping detail, or professionally remove backgrounds from product photos. This MCP lets you handle the entire visual media pipeline right through conversation.

A+ Quality Score 100/100

generative-art | image-generation | upscaling | background-removal | visual-media | computer-vision

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

Stability AI MCP

10 tools available
Cloud-hosted on Vinkius

This connector brings industry-leading generative visual capabilities into your existing LLM workflows. Instead of jumping between multiple design tools and juggling complex APIs, you simply talk to your agent. You can ask it to generate a full mockup for an ad campaign from scratch using only text prompts. Need better quality? Just tell it to upscale the image later. Or maybe you've got product photos that need cleaning up; this MCP handles background removal instantly. It's all about iteration: generating drafts, refining details, and improving resolution—all in one place. If your team uses Vinkius for other services, adding Stability AI gives you a single point of access to professional visual creation.

Core Capabilities

01 — Generate Images from Text

Create entirely new images using advanced models based on detailed written instructions.

02 — Improve Image Resolution

Increase the size and detail of existing graphics without losing structural quality or fidelity.

03 — Remove Backgrounds

Automatically clean product photos by isolating the subject and removing its background.

04 — Edit Specific Image Areas

Replace isolated parts of a photo, like blemishes or missing elements, while keeping the surrounding area consistent.

One Click on Vinkius — From Prompt to Execution

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

01 Connect your API key to this MCP within your agent's settings.

02 Tell your agent exactly what you need. For example: 'Generate a futuristic bike mockup and then remove its background.'

03 Your agent executes the necessary steps using multiple tools, delivering the completed, refined visual asset.

The bottom line is that instead of running three different processes in separate software, your AI client manages the entire creative workflow for you.

Built For

This MCP solves problems for anyone who touches visual media professionally. It's for marketing managers who can't wait on designers, e-commerce owners needing consistent product shots in bulk, and concept artists who need to rapidly prototype ideas.

E-Commerce Manager

You need to standardize thousands of product photos daily. You use this MCP to automatically remove backgrounds and generate cohesive catalog imagery for your site.

Digital Marketing Specialist

Your campaign requires 50 ad variations testing different styles. You ask the agent to generate diverse mockups instantly, using various generation models to test ideas fast.

Creative Operations Lead

You need to take rough sketches or low-res concepts and turn them into high-fidelity presentations for stakeholders without manually editing every detail.

What Changes When You Connect

01 You stop worrying about which generation model to use. You just ask for the result you want—be it `generate_core_v2` speed or `generate_ultra_v2` quality—and your agent handles the selection.

02 Need to adapt an existing mockup? Use `image_to_image_v1` to change the style of a generated image, keeping its composition while applying a new look. It's powerful iteration in one prompt.

03 Product photography just got easier. The `remove_background` tool cleans up product shots instantly, standardizing your catalog images without Photoshop.

04 Low-resolution assets don't have to be discarded. Simply run `upscale_image` and get a high-detail version that can actually be used in print or digital ads.

05 The system handles complex edits with tools like `inpaint_image`, letting you fix small errors or add missing details to an image without having to mask off the area first.

Real-World Applications

Updating a product catalog in bulk

An e-commerce manager needs 50 shirts photographed against different backdrops. Instead of hiring a photo studio, they instruct their agent to run the `remove_background` tool on all photos, then ask the agent to generate new background images and composite them automatically.

Creating rapid ad campaigns

A digital marketing specialist needs 20 different futuristic car mockups for A/B testing. They use `generate_sd35` with a specific prompt, asking the agent to cycle through multiple models and generate variations quickly.

Reviving old artwork

A creative director has an old concept sketch that is too low-res. They use `upscale_image` on the sketch, then refine specific details using `inpaint_image` to turn it into a high-resolution pitch piece.

Developing complex mockups

A product designer needs an image of a smart speaker in a modern living room. They prompt the agent to generate the initial mockup using `generate_core_v2`, and then follow up by asking it to adjust the lighting or remove extraneous objects.

Patterns to Avoid

Treating generation like a single step

✗ AVOID

Trying to generate an image, then manually downloading it, opening Photoshop, and running background removal separately.

✓ INSTEAD

Instead, tell your agent: 'Generate the product mockup, and immediately use `remove_background` on the result.' This keeps the entire process contained within one conversation.

Forgetting about different models

✗ AVOID

Using a basic generation tool when you need final-grade marketing assets, resulting in blurry or inconsistent output.

✓ INSTEAD

When generating for print or client pitches, specify using the high-quality `generate_ultra_v2` model to ensure maximum fidelity.

Assuming simple editing is enough

✗ AVOID

Simply asking the agent to 'edit the photo' without specifying what needs changing.

✓ INSTEAD

Be specific. Use `inpaint_image` and tell it exactly which area to change (e.g., 'change the color of the chair in this spot') rather than using a general edit command.

The Right Fit

Use this MCP if your core need is visual asset creation, refinement, or systematic cleanup. You're building a pipeline: generate -> refine -> clean. Don't use it if you just need to analyze existing images (that requires an image classification tool). Also, don't use it if you only need simple text-to-text content; that's for language models alone. However, this MCP is ideal when your workflow demands multiple passes: first generating the concept (`generate_sd35`), then

making minor changes to a specific object using `inpaint_image`, and finally increasing the size with `upscale_image`. If you are only doing one thing—say, just background removal—you could use a specialized category tool, but this MCP gives you maximum control over the entire process.

Manual visual asset pipelines waste hours of time.

Right now, creating a single marketing ad requires juggling five different applications. You generate the image in tool A, download it, open it in Photoshop B to cut out the background, then use an online service C to upscale it for print, and finally compile all those files into your campaign tracker D.

With this MCP, you just talk to your agent. You describe the final outcome—'I need a high-res product shot with a clean white background.' The agent handles the entire sequence: generating the initial image, isolating the subject using `remove_background`, and ensuring the output is ready for use.

Control your visual assets with Stability AI MCP

The biggest time sinks disappear. You no longer waste minutes opening, saving, and renaming files across multiple desktop programs. The agent keeps the entire context of the image in memory.

Now, you simply define the desired state of the image, and your MCP delivers it. It changes everything from a multi-hour manual process to a single conversation turn.

Stability AI MCP: 10 Tools for Visual Media

These tools allow you to control every step of the visual media creation process, from initial text prompting to final quality enhancement.

#	TOOL	DESCRIPTION
01	`generate_core_v2`	Creates an image using the fast and quality-optimized Stable Image Core model.
02	`generate_sd35`	Generates images using Stable Diffusion 3.5, allowing you to choose between three specific models (large, turbo, or medium).
03	`generate_ultra_v2`	Produces high-end photorealistic images best suited for final production assets.
04	`get_credit_balance`	Checks your current usage quota and available Stable AI credits.
05	`image_to_image_v1`	Modifies an existing picture by transforming it based on a new text prompt.
06	`inpaint_image`	Changes specific, small regions within an image using a prompt while keeping the rest of the photo intact.
07	`list_engines`	Lists all available engine IDs you can use for generating images with v1 tools.
08	`remove_background`	Automatically strips the background from any given photograph.
09	`text_to_image_v1`	Generates an image from scratch using a text prompt and specified size parameters via v1 engines.
10	`upscale_image`	Increases the resolution of an image while mathematically preserving its visual details.

See It in Action

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

U **Generate a wide format concept visual depicting a sleek, futuristic electric bike stationed alongside a minimalist architectural wall structure.**

Formatting inputs processing core parameters executing the request systematically through the established model engine...

- **Processing Protocol**: Core Generation (`generate_image`)
- **Render Resultant Instance**: [Image Received via API Data Successfully Base64]

Do you wish to iterate computationally adjusting resolution density with `upscale_image`, or are requirements sufficiently fulfilled?

U **Upscale this low-resolution image of a landscape without losing structural fidelity.**

Command processed. Executing `upscale_image` logic... Done. Resolution enhanced successfully.

U **Remove the background from this product photography.**

Executing `remove_background` routine... Done. Background eradicated cleanly.

Frequently Asked Questions

01 **How do I use Stability AI with my existing images?**

You can transform or enhance old photos using `image_to_image_v1`. This tool lets you apply a new style to an image while keeping the original structure intact.

02 **Is the generate_ultra_v2 model best for everything?**

No. While `generate_ultra_v2` is great for final assets, if speed is your main concern, you might prefer using models like those in the `generate_core_v2` tool.

03 **Can I generate images at a specific size?**

Yes. When generating from text prompts via v1 engines, you can provide explicit width and height parameters to control the output dimensions.

04 **What if my product image has complex elements in the background that need fixing?**

Use `inpaint_image`. This tool lets your agent surgically target a specific area of the photo, allowing you to clean up or replace objects without affecting the surrounding detail.

05 **Do I have to pay for every generation with Stability AI MCP?**

No. You can check your current consumption and remaining budget by calling the `get_credit_balance` tool before starting a large batch of operations.

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 `"stability-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

Stability 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 Stability 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	Stability AI MCP
Server ID	`019d760c-b536-7123-9a21-8fa6ddaa1de7`
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/stability-ai.