MCP SERVER | NO CODE | CLOUD HOSTED

# Creative Residency Decision Support MCP

## Decide which residencies fit your life and goals.

Creative Residency Decision Support MCP helps you weigh residency opportunities against your personal constraints. It analyzes host terms, costs, and accessibility needs to tell you whether to apply, defer, or pass on an offer. You can use it to build application strategies, draft questions for organizers, and plan how to balance a residency with your existing work.

A+ Quality Score 100/100

residency | artist | planning | decision-making | creative-process

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026
vinkius.com/en/ai-agent-connect/creative-residency-decision-support

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.

# Creative Residency Decision Support MCP

5 tools available
Cloud-hosted on Vinkius

Choosing a residency is rarely just about the art. You have to weigh travel costs, studio availability, and scheduling conflicts against your current workload and physical needs. This MCP acts as a logic engine for those decisions. Instead of guessing if an opportunity makes sense, you feed the residency details into your AI client and get a clear breakdown of the pros and cons.

You can use it to vet specific offers or prepare for ones you already want. If a residency looks good but the timing is off, the MCP helps you figure out how to ask the host for a different window or how to manage your current projects so you can actually go. It moves you from vague interest to a concrete plan, covering everything from the initial application steps to a backup plan if things don't go as expected.

## Core Capabilities

**01 — Decision Logic**

Your agent uses this to weigh residency terms against your specific constraints.

**02 — Application Planning**

Your agent generates actionable steps to complete a residency application.

**03 — Risk Management**

Your agent creates contingency plans to handle potential scheduling or financial issues.

**04 — Communication Drafting**

Your agent writes targeted questions for residency hosts to resolve dudas.

**05 — Schedule Balancing**

Your agent outlines how to manage a residency alongside your current commitments.

## One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/creative-residency-decision-support** — connect your AI agent in three steps.

01 Connect your AI client to the Vinkius-hosted MCP.

02 Provide the residency details and your personal constraints to your agent.

03 Run the evaluation to get an Apply, Decline, or Defer recommendation.

04 Use the resulting strategy or commitment plan to take action.

Connect the MCP to your preferred AI client to start analyzing opportunities.

## Built For

This MCP is built for artists and creative professionals who need to manage complex schedules and opportunities.

**Visual Artists**

They hand off residency terms and studio requirements to the AI to check for fit.

**Writers**

They use the tool to plan application timelines and manage writing commitments.

**Performing Artists**

They use it to balance touring or performance schedules with residency offers.

## What Changes When You Connect

01 Reduces decision fatigue by weighing costs and accessibility automatically.

02 Provides clear action steps for residency applications.

| | |
|---|---|
| 03 | Helps manage existing work commitments during residency periods. |
| 04 | Prepares backup plans for high-risk opportunities. |

## Real-World Applications

**Vetting New Offers**

Compare a new residency's cost and location against your current budget and availability.

**Resolving Scheduling Conflicts**

Determine if a residency is worth deferring due to a conflicting work project.

**Preparing Applications**

Get a checklist of portfolio and statement requirements for a specific program.

**Host Communication**

Generate specific questions about studio accessibility or facility needs.

## Patterns to Avoid

### Vague decision making

AVOID

You tell your agent you want to go to a residency in France but don't provide any details about your budget or current projects.

✓ INSTEAD

Provide specific constraints like travel costs and existing deadlines, then use evaluate_residency_options to get a clear recommendation.

### Ignoring potential conflicts

AVOID

You accept a residency offer without checking if your current studio commitments or travel expenses will make it impossible to attend.

✓ INSTEAD

Use create_commitment_plan to map out how your existing obligations will interact with the new residency schedule.

### Assuming all details are clear

AVOID

You assume the residency provides all the equipment you need and skip asking the host about studio accessibility.

✓ INSTEAD

Use formulate_host_inquiries to generate a list of specific questions for the organizers before you commit.

## The Right Fit

Connect this MCP if you are a creative professional managing multiple overlapping schedules and financial constraints. Do not use it if you have a single, predictable workload with no travel or studio requirements. The decision depends on whether your opportunities require weighing logistical pros and cons against your current life.

# 5 Tools for Residency Planning and Decision Support

These tools handle everything from initial vetting to building a full contingency plan for your residency.

| # | TOOL | DESCRIPTION |
|---|---|---|
| 01 | `draft_contingency_arrangement` | This tool builds a backup plan based on your primary decision and any identified risks. |
| 02 | `evaluate_residency_options` | This tool determines if you should apply, decline, or defer a specific residency. |
| 03 | `generate_application_strategy` | This tool outlines the specific steps you need to take to pursue a residency you've decided to apply for or defer. |
| 04 | `create_commitment_plan` | This tool maps out how you will balance a residency with your existing work and life obligations. |
| 05 | `formulate_host_inquiries` | This tool generates a list of questions to send to residency organizers to clear up uncertainties. |

# See It in Action

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

**U** Should I apply for the Summer Arts Residency? It costs $2000 and provides a studio, but I have a work commitment in July.

The decision is Defer. While the studio meets your goals, the July work commitment creates a scheduling conflict that needs to be resolved with the host first.

**U** I want to apply for the Mountain Retreat. Help me plan my application.

To pursue the Mountain Retreat, you should first finalize your portfolio, draft your project statement focusing on nature-inspired work, and prepare your letters of recommendation.

**U** What questions should I ask the host about accessibility?

You should ask: 'Does the studio space have step-free access?' and 'Are there specific height-adjustable workstations available in the facility?'

# Frequently Asked Questions

01 **What AI clients can I use with this MCP?**

You can use this MCP with any compatible client, including Claude, Cursor, Windsurf, and VS Code.

02 **How does the MCP help with residency decisions?**

It weighs the host's terms, such as cost and accessibility, against your specific goals and constraints to suggest whether to apply, decline, or defer.

03 **Can it help me write emails to residency hosts?**

Yes, the tool can generate specific inquiries to help you resolve uncertainties with organizers.

04 **Does this MCP handle my scheduling?**

It creates a commitment plan that outlines how to balance a residency with your existing life and work obligations.

05 **Where is the MCP hosted?**

Vinkius hosts and manages the MCP, so you don't need to run it locally or manage credentials yourself.

06 **How does the tool decide to 'Decline' a residency?**

A residency is automatically declined if it violates any of your mandatory terms, such as exceeding your maximum cost threshold or failing to meet specific accessibility requirements.

07 **What is the difference between 'Defer' and 'Decline'?**

Decline means the residency is a non-starter due to hard constraints. Defer means the residency is interesting, but you need more information (like specific accessibility details) before you can commit.

08 **Can I use this to manage my schedule?**

Yes, you can use `create_commitment_plan` to outline how you will balance your existing obligations with the residency timeframe.

# 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 `"creative-residency-decision-support": { "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

# Creative Residency Decision Support 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 Creative Residency Decision Support. 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 | September 2026 |
| MCP Server | Creative Residency Decision Support MCP |
| Server ID | `01a0dabe-b512-71d9-9973-1c7db234f5f2` |
| 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/en/ai-agent-connect/creative-residency-decision-support.

© 2026 Vinkius · vinkius.com · support@vinkius.com