MCP SERVER | NO CODE | CLOUD HOSTED

Civic Hearing Participation Planner MCP

Organize your testimony and logistics for official civic hearings.

Civic Hearing Participation Planner MCP gives you a structured way to handle the chaos of formal civic proceedings. It helps you map out preparation milestones, manage attendee registration, draft timed speeches, and coordinate your follow-up actions after the hearing ends. You can use it directly within your AI client to turn raw hearing data into a concrete action plan.

A+ Quality Score 100/100

civic | hearing | testimony | planning | logistics

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026
vinkius.com/en/ai-agent-connect/civic-hearing-participation-planner

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.

Civic Hearing Participation Planner MCP

4 tools available
Cloud-hosted on Vinkius

Public hearings and civic meetings are often rigid and high-stakes. If you miss a registration deadline or run over your allotted speaking time, your input might be ignored entirely. This MCP changes that by giving your AI client the specific tools needed to manage the logistics of civic engagement.

You can use it to build a step-by-step timeline so you aren't scrambling the night before a meeting. It handles the heavy lifting of organizing your thoughts into a timed speech, ensuring you hit your key points without getting cut off by the chair. Beyond the hearing itself, it helps you stay organized by planning for registration and managing what happens once the session concludes. Instead of staring at a blank page or a messy calendar, you can use your agent to build a professional, timed, and organized presence at any local government or civic proceeding.

Core Capabilities

01 — Timeline Generation

Your agent creates a sequence of deadlines and milestones to ensure you are ready by the hearing date.

02 — Speech Structuring

Your AI client builds timed outlines to prevent you from exceeding your allotted speaking time.

03 — Logistics Management

Your agent handles registration and attendance planning for participants.

04 — Post-Hearing Coordination

Your agent generates specific action items to follow up with committees after the meeting.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/civic-hearing-participation-planner** — connect your AI agent in three steps.

01 Connect your AI client to the Vinkius platform.

02 Provide your hearing date and specific requirements to your agent.

03 Use the tools to generate timelines, speech outlines, or attendance plans.

04 Execute your plan and use the follow-up tool to close out the process.

Connect your preferred client to Vinkius to access these tools immediately.

Built For

This MCP is built for anyone who needs to present formal arguments or manage attendance at government or civic meetings.

Community Advocates

They hand off meeting dates and goal descriptions to have the AI build a full preparation and testimony plan.

Local Government Staff

They use the tool to manage registration logistics and follow-up tasks for public sessions.

Legal Professionals

They use the tool to structure timed testimony and manage post-hearing documentation tasks.

What Changes When You Connect

01 Prevents being cut off during testimony by using timed speech outlines.

02 Reduces missed deadlines through automated preparation timelines.

03	Ensures post-hearing actions are completed via structured follow-up plans.
04	Organizes participant logistics through centralized attendance planning.

Real-World Applications

Zoning Board Testimony

An individual uses the tool to structure a 3-minute speech regarding local land use changes.

Community Group Mobilization

A group uses the attendance tool to manage registration for a large number of interested citizens.

Environmental Impact Advocacy

An advocate uses the timeline tool to prepare documents and travel before a scheduled hearing.

Post-Meeting Follow-up

A participant uses the follow-up tool to ensure their submitted evidence is officially reviewed by a committee.

Patterns to Avoid

Ignoring time limits

AVOID

You write a long, rambling speech and hope the board chair lets you finish. You end up getting cut off before you reach your main point.

✓ INSTEAD

Use `structure_testimony` to turn your thoughts into a timed outline that fits within the official limits.

Missing registration windows

AVOID

You show up to a zoning meeting only to find out the deadline for public comment passed two days ago.

✓ INSTEAD

Run `generate_preparation_timeline` to map out every deadline and milestone before the hearing date.

Neglecting post-meeting follow-up

AVOID

You deliver your testimony and assume the committee will automatically act on it. Nothing happens because you didn't track the next steps.

✓ INSTEAD

Use `generate_follow_up_plan` to create a specific list of actions to ensure your input gets processed.

The Right Fit

Connect this MCP if you regularly present formal arguments or manage attendance for civic meetings. Do not connect it if you are just a casual observer with no intention of testifying or organizing participants. The deciding factor is whether you need to manage strict deadlines and timed speaking slots.

Civic Hearing Participation Planner Tools

These tools handle the logistics, timing, and follow-up required to participate effectively in public government proceedings.

#	TOOL	DESCRIPTION
01	`create_attendance_plan`	This tool manages the logistics and registration requirements for your hearing. It helps you keep track of who needs to be there and what registration steps are required.
02	`generate_follow_up_plan`	This tool builds an action plan for the period after a hearing concludes. It helps you track necessary next steps to ensure your input is actually processed.
03	`generate_preparation_timeline`	This tool creates a structured timeline of milestones leading up to your hearing date. It helps you stay on track with deadlines and preparation tasks.
04	`structure_testimony`	This tool organizes your speech into a timed outline. It ensures your testimony remains effective and stays within the official time limits.

See It in Action

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

U I have a hearing on 2025-05-15. I need to travel and my materials will be ready on 2025-05-01.

Your preparation timeline includes milestones for material finalization on May 1st, travel booking, and registration deadlines leading up to May 15th.

U Help me organize my speech. I have 3 minutes (180 seconds) to discuss zoning laws and environmental impact using the provided impact report.

Your testimony outline is structured into an Introduction, Argument regarding zoning and environmental impact, and a Conclusion, with a total estimated time of 160 seconds to ensure you stay within your 180-second limit.

U What should I do after my hearing on 2024-10-10? I want to ensure the committee reviews my submitted documents.

Your follow-up plan includes high-priority tasks for documenting the hearing outcome and communicating with the committee to confirm receipt of your documents.

Frequently Asked Questions

01 **How does this MCP help with timed speeches?**

The structure_testimony tool organizes your points into a timed outline. This helps you stay within the specific time limits set by the hearing chair.

02 **Can I use this with Claude or Cursor?**

Yes. You can connect this MCP to any MCP-compatible client like Claude, Cursor, or Windsurf through the Vinkius platform.

03 **Does this help with post-hearing tasks?**

Yes. The generate_follow_up_plan tool creates a specific list of actions to take after the hearing is over.

04 **What if I am managing a large group of people?**

You can use the create_attendance_plan tool to manage registration and the logistics of your participants.

05 **Do I need to host this myself?**

No. Vinkius hosts and manages the MCP for you. You just connect your client and start using the tools.

06 **How does this tool help with speaking limits?**

The `structure_testimony` tool automatically calculates a buffer to ensure your speech stays within the official speaking limits.

07 **Can I use this for informal meetings?**

While designed for formal hearings, the tools can be adapted for any meeting requiring structured preparation and follow-up.

08 **Does it handle travel logistics?**

Yes, `generate_preparation_timeline` includes specific milestones for travel if you indicate that travel is required.

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 `"civic-hearing-participation-planner": { "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

Civic Hearing Participation Planner 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 Civic Hearing Participation Planner. 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	Civic Hearing Participation Planner MCP
Server ID	`01a0d6a1-e234-72f2-949f-f059d7c120a1`
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/civic-hearing-participation-planner.