MCP SERVER | NO CODE | CLOUD HOSTED

Local Event Attendance Planner MCP

Coordinate your entire day without missing a single event.

Local Event Attendance Planner MCP handles the math of your social calendar. It calculates travel times, manages your budget, and finds the best sequence of events so you can spend more time at the venue and less time worrying about how to get there or if you can afford the next ticket.

A+ Quality Score 100/100

itinerary | travel | events | optimization | planning

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026
vinkius.com/en/ai-agent-connect/local-event-attendance-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.

Local Event Attendance Planner MCP

4 tools available
Cloud-hosted on Vinkius

You have a list of events you want to attend, but the logistics are a mess. This MCP acts as your personal scheduling engine. Instead of manually checking maps and ticket prices, you give your constraints to your AI client and let it do the heavy lifting. It looks at your budget, your need for wheelchair accessibility, and how much buffer time you need between stops to build a logical plan. If your first choice is too expensive or the timing is too tight, it finds a different way to make it work. It doesn't just give you a list; it gives you a precise timeline for when to leave your house and exactly what you need to buy to make the day happen.

Core Capabilities

01 — Constraint-based scheduling

The AI uses your budget and accessibility needs to pick the right events.

02 — Travel buffer calculation

The AI adds necessary time between stops to prevent you from being late.

03 — Budget management

The AI checks total costs against your limit before finalizing a plan.

04 — Logistics identification

The AI lists the specific tickets and transport you need to purchase.

05 — Timeline generation

The AI creates specific departure times for every leg of your trip.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/local-event-attendance-planner** — connect your AI agent in three steps.

01 Connect your MCP-compatible client to the Vinkius catalog.

02 Provide your list of events, budget, and travel preferences to your AI client.

03 The AI uses the MCP to calculate the best itinerary and departure times.

04 Review the booking actions to see what tickets you need to buy.

05 Follow the generated timeline to complete your day.

Connect your client to Vinkius and start planning immediately.

Built For

People who want to maximize their time at local venues without the mental load of planning.

Event enthusiasts

They hand off complex lists of shows and festivals to the AI to find the best way to attend them all.

Accessibility planners

They use the tool to ensure every stop in a sequence meets specific physical requirements.

Budget-conscious travelers

They let the AI filter event combinations to stay under a strict spending limit.

What Changes When You Connect

01 Eliminates manual math for travel buffers and transit times.

02 Prevents overspending by checking costs against your budget.

03	Reduces the risk of missing events due to poor timing.
04	Provides clear instructions on what to buy to secure your spot.

Real-World Applications

Festival Hopping

Coordinate multiple music performances while staying within a set budget.

Accessible Sightseeing

Build a day of museum visits that strictly adhere to wheelchair access needs.

Tight Schedule Management

Use precise departure times to move between multiple city locations without being late.

Budget Re-routing

Find a cheaper set of events when your primary plan is too expensive.

Patterns to Avoid

Vague constraint input

AVOID

Telling your agent you want to go to a few shows today without mentioning a budget or mobility needs.

✓ INSTEAD

Provide specific limits on spending and accessibility requirements. Use plan_itinerary to build a sequence that actually fits your real-world constraints.

Ignoring the timeline

AVOID

Assuming you have enough time to get from one venue to another without checking the transit gaps.

✓ INSTEAD

Always check your specific departure times. Use get_departure_plan to see exactly when you need to walk out the door to stay on schedule.

Missing the purchase list

AVOID

Planning a full day of festivals but forgetting that you need specific transit passes or timed entry tickets.

✓ INSTEAD

Run get_booking_actions after your plan is set. It identifies the exact tickets and transport you need to buy so you aren't stuck at the gate.

The Right Fit

Connect this MCP if you frequently manage schedules involving multiple locations, tight budgets, or specific accessibility needs. Do not connect it if you are only attending a single event at a single location. The decision rests on whether you need to solve for transit buffers and cost constraints across a sequence of stops.

Local Event Attendance Planner Tools

These tools handle the math and logistics for your event schedules. They turn a list of interests into a concrete, actionable timeline.

#	TOOL	DESCRIPTION
01	`get_alternate_plans`	This tool suggests different event combinations if your first itinerary doesn't fit your needs.
02	`get_booking_actions`	This tool identifies the specific tickets or transport passes you need to buy to execute your plan.
03	`get_departure_plan`	This tool provides a granular timeline so you know exactly when to leave to arrive on time.
04	`plan_itinerary`	This tool builds the optimal sequence of events based on your specific budget and travel requirements.

See It in Action

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

U Plan an itinerary for 2 people with wheelchair access. Events: Concert at 6pm ($50), Museum at 2pm ($20). Budget: $100. Travel buffer: 30 mins.

Your itinerary includes the Museum at 2:00 PM and the Concert at 6:00 PM. Total cost is $70, which is within your $100 budget. Both locations are wheelchair accessible.

U What do I need to book for my planned itinerary?

You need to buy a ticket for the Jazz Festival and reserve a ride via the public transit service.

U Give me a departure timeline for my museum visit starting at 2pm.

Leave your current location by 1:15 PM to arrive at the Museum by 2:00 PM.

Frequently Asked Questions

01 **How does this MCP handle travel time?**

It uses a travel buffer to calculate the time needed between events, ensuring you arrive on time.

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.

03 **What if my planned itinerary is too expensive?**

You can ask the AI to find alternate plans that fit within your specific budget.

04 **Does it help with accessibility needs?**

Yes, you can include accessibility requirements in your prompt, and the tool will account for them when planning.

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

No, Vinkius hosts and manages the MCP for you, so it is ready to use once you connect.

06 **How does the tool handle travel time between events?**

The engine applies a mandatory travel buffer between events to ensure you have enough time to move between locations based on your selected transport mode.

07 **Can I plan for group attendance with specific needs?**

Yes, you can provide a list of companions and their accessibility requirements. The engine ensures all selected events and transport modes are accessible to everyone in the group.

08 **What happens if my budget is too low for the preferred events?**

You can use `get_alternate_plans` to find different itinerary combinations that prioritize your needs while staying strictly within your budget limits.

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 `"local-event-attendance-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

Local Event Attendance 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 Local Event Attendance 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	Local Event Attendance Planner MCP
Server ID	`01a0d5f0-50be-7265-ba2c-29b234af8b30`
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/local-event-attendance-planner.