MCP SERVER | NO CODE | CLOUD HOSTED

TomTom Parking Availability MCP

Audit real-time parking spots from any conversation.

TomTom Parking Availability provides real-time data to search for parking spots anywhere in an urban area. This MCP lets your AI client audit locations, check facility availability using specific IDs, and verify detailed coordinates instantly. It turns complex logistics research into a simple conversation, giving you precise, localized mobility intelligence without needing any dedicated navigation app.

A+ Quality Score 100/100

parking-availability | urban-mobility | location-services | real-time-data | map-data | infrastructure-auditing

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

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.

TomTom Parking Availability MCP

3 tools available
Cloud-hosted on Vinkius

Managing city-scale logistics or running regional audits shouldn't feel like juggling ten different apps. With TomTom Parking Availability connected via Vinkius, your agent handles the heavy lifting. You can ask it to find thousands of parking spots near a specific intersection, pulling back detailed metadata including facility names and addresses. It doesn't just give you an estimate; it audits the exact geographic coordinates for those sites, helping you understand local distribution instantly. Need to know how full a garage is? Give your agent a specific ID, and it pulls the latest availability data. This means whether you're managing a large fleet or doing urban planning research, your AI client acts as a real-time mobility consultant that keeps your data accurate and actionable.

Core Capabilities

01 — Search for parking spots by location

The agent finds multiple potential parking facilities near any address you specify.

02 — Audit specific facility availability

You check the current occupancy status and full details using a known parking ID.

03 — Verify coordinates for sites

The agent retrieves precise geographic data, allowing you to map out regional distribution.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/tomtom-parking-availability** — connect your AI agent in three steps.

01 Subscribe to this MCP and provide your TomTom API Key.

02 Connect the service to your preferred AI client (Claude, Cursor, etc.) through Vinkius.

03 Ask your agent a natural language question, like 'Find parking spots near downtown Denver' or 'What is ID 54321 availability?'

The bottom line is you tell your agent what data you need in plain English, and it executes the complex API calls behind the scenes.

Built For

This MCP is essential for anyone whose job involves physical location planning or managing mobile assets. If you're an operations lead who gets frustrated by inaccurate manual data collection, or a logistics planner needing instant spot checks across multiple city zones, this tool saves hours of tedious searching.

Logistics Planner

You use the MCP to monitor parking availability and retrieve location metadata directly within your planning workflow.

Fleet Manager

You verify multiple parking addresses across different regions, auditing regional distribution without manual map searches.

Urban Researcher

You perform rapid audits of Point-of-Interest (POI) data and identify relevant mobility markers using natural language prompts.

What Changes When You Connect

01 Instantly find multiple options: By using the `search_parking_spots` tool, your agent finds thousands of nearby facilities and returns detailed metadata like facility names and addresses, eliminating manual map searching.

02 Know availability now: The `get_parking_details` function allows you to check the current occupancy status for a specific parking ID, giving immediate insight into resource capacity.

03 Pinpoint locations accurately: You can use coordinates to verify precise geographic data. This capability helps regional managers understand exact local distribution patterns efficiently.

04 Streamline auditing workflows: Instead of jumping between multiple mapping services, your agent consolidates all necessary location and availability checks in one conversation pane.

05 Maintain operational integrity: The `check_api_status` tool ensures that your mobility research workflow is always functional before you rely on its output.

Real-World Applications

Checking a new branch location's parking capacity

The operations lead needs to know if the proposed retail site has adequate nearby parking. They prompt their agent: 'Search for parking spots near 123 Main St.' The agent uses `search_parking_spots` and returns multiple options, allowing the lead to immediately verify location suitability.

Assessing regional fleet density

A fleet manager needs to compare parking capacity across three different zip codes. They ask their agent to audit coordinates in a specific zone. The agent uses coordinate lookups and `search_parking_spots` to build an instant, comprehensive map of available assets.

Troubleshooting logistics bottlenecks

The dispatcher needs immediate status on a known garage ID because a delivery is delayed. They ask their agent: 'What's the availability for parking ID 9987?' The agent uses `get_parking_details` to return the current occupancy rate, solving the delay instantly.

Planning a major urban development project

The city planner needs to know all possible POI parking sites within a 0.5-mile radius of a new transit hub. They prompt their agent, and it uses coordinate data retrieval alongside `search_parking_spots` to compile an exhaustive list for the proposal.

Patterns to Avoid

Treating location data as static

✗ AVOID

Copy-pasting a general address into a map and assuming it shows live parking availability or capacity.

✓ INSTEAD

Use the `search_parking_spots` tool with your agent. This searches for facilities *near* that address, providing multiple POIs, not just one static pin.

Forgetting to verify service uptime

✗ AVOID

Writing a long workflow assuming location data is available when the underlying API connection might be down.

✓ INSTEAD

Always start by calling `check_api_status`. This confirms the entire mobility research pipeline is operational before you run costly searches.

Overlooking specific facility IDs

✗ AVOID

Asking for 'parking downtown' and getting a list, but having no way to check if any single site is full or nearly full.

✓ INSTEAD

Use `get_parking_details` with the specific ID provided by your agent. This gives you real-time occupancy data, which is much more valuable than just knowing the name.

The Right Fit

You should use this MCP if your core problem involves verifying physical location details or needing to know how many spots are available right now. If you need to audit a region's capacity or check if an asset is operational, this is the tool for you.

Don't use this if your task is purely theoretical—like mapping out potential zoning changes that have no current infrastructure data. Similarly, don't try to calculate hypothetical traffic flow based on historical averages; this MCP focuses on real-time location intelligence and availability. If all you need is a simple 'directions to X,' then standard map services suffice. But if your workflow requires querying specific facility IDs or running multi-point searches across a city grid, you need the precision of `search_parking_spots` combined with the real-time checks from `get_parking_details`. It's built for operational auditing.

The headache of manual site audits is exhausting.

Today, if you need to audit a new multi-site location, you have to open three different tabs: one for Google Maps coordinates, another for general POI data, and then maybe a third site's website just to see if they even mention parking. You copy an address here, paste it there, wait for the map to load, realize it only gives you a pin, and then start guessing about availability.

With this MCP, your agent handles the whole thing in one prompt. Instead of clicking through multiple services, you ask for location details, available spots, and even specific IDs—and the answer comes back compiled into one conversational response. You get actionable intelligence instead of a mess of links.

Get real-time parking data with TomTom Parking Availability.

You eliminate the need to manually check if an address is accurate, then separately verify its coordinates, and finally make a second trip just to

Your agent provides immediate, consolidated data on parking availability and location metadata—it's

find out if that location even has available spots. All these steps collapse into one query.

real-time operational intelligence you can actually trust.

TomTom Parking Availability: 3 Tools

These tools let you programmatically search for potential parking facilities, retrieve real-time availability data using specific IDs, and audit precise geographic coordinates.

#	TOOL	DESCRIPTION
01	`check_api_status`	Checks if the TomTom Parking service is currently running and available for use.
02	`get_parking_details`	Retrieves full details and current availability status for a specific parking facility ID.
03	`search_parking_spots`	Searches for multiple potential parking facilities near a given location or set of coordinates.

See It in Action

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

U Search for parking spots in 'San Francisco' using TomTom.

I've retrieved several parking facilities in San Francisco! Notable options include the 'Fifth & Mission Garage' and 'Union Square Garage'. Would you like the full addresses or coordinate metadata for these locations?

U What is the availability for parking ID '12345'?

I've retrieved the details for parking ID 12345! It is currently identified as having 'High' occupancy. I can assist you with the exact address and location metadata if you'd like to identify nearby alternatives.

U Find parking near latitude 37.7749 and longitude -122.4194.

I've identified the nearest parking facilities via your coordinates! These include several garages within a 0.5-mile radius. I can provide the specific IDs and names for each of these sites to assist in further mobilty audits.

Frequently Asked Questions

01 **How does TomTom Parking Availability find spots near a specific address?**

The `search_parking_spots` tool executes a search around your provided location. It returns multiple nearby facilities and crucial metadata, allowing you to see several options instantly.

02 **Can I check the current capacity of a parking garage using TomTom Parking Availability?**

Yes, use `get_parking_details`. You provide a specific parking ID, and the MCP returns real-time occupancy data for that exact facility.

03 **Is this service just for city planning or can I use it for logistics?**

It's built for both. Logistics planners use it to verify regional distribution and capacity, while researchers use it for detailed urban mapping and auditing.

04 **What if my search area is too large for TomTom Parking Availability?**

The agent handles the scope; you just need to provide a clear target location or coordinate range. The tool pulls data relevant to that defined boundary.

05 **Do I need a separate API key if I use TomTom Parking Availability?**

Yes, you must subscribe and enter your unique TomTom API Key into the MCP setup before it can run any tools or queries.

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 `"tomtom-parking-availability": { "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

TomTom Parking Availability 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 TomTom Parking Availability. 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	TomTom Parking Availability MCP
Server ID	`019d8490-35ef-71ec-9002-f317634d2844`
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/tomtom-parking-availability.