MCP SERVER | NO CODE | CLOUD HOSTED

TfL Transit MCP

Know London’s Live Transport Status Instantly

Transport for London MCP gives your AI agent instant access to real-time data across all of London's transport systems. Plan multi-modal journeys, check live tube line status, find bus arrivals at any stop, or track major road disruptions—all through natural conversation without needing an API key.

A+ Quality Score 100/100

public-transport | transit-data | real-time-updates | journey-planning | live-status | urban-mobility

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

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.

Transport for London MCP

11 tools available
Cloud-hosted on Vinkius

Need to know if the Victoria Line is running good today? Or maybe you just need directions from King's Cross to Heathrow that account for a potential delay on the Jubilee line. This MCP connects your AI agent directly to London's transit data, giving it instant awareness of what's happening in real time. You can ask complex questions—like 'What is the fastest route to Oxford Circus right now?'—and get detailed directions that consider walking distance, expected cost, and connection changes. It covers everything from bus arrivals at a specific stop to finding available bikes at a nearby docking station. Because this MCP sits on Vinkius, you don't have to worry about managing credentials or integrating multiple data feeds; your agent just knows how to talk to the entire system.

Core Capabilities

01 — Plan Multi-Modal Routes

Generates step-by-step directions between any two London locations, providing estimated duration and fare cost.

02 — Check Tube Line Status

Provides the real-time service status for every Underground line, indicating if it's running smoothly or experiencing delays.

03 — Get Live Bus Arrivals

Predicts exactly when buses will arrive at a specific stop point in London.

04 — Monitor Road Conditions

Reports on major road congestion levels or details about current traffic disruptions across the city.

05 — Find Cycling Stations

Locates nearby bike docking stations and tells you how many bikes are available at each one.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/transport-for-london** — connect your AI agent in three steps.

01 Connect your preferred AI client to this MCP via Vinkius. This gives the agent access to all live London transit data.

02 Ask a natural language question, like 'What's the status of the Northern line?' or 'How do I get from Waterloo to Tower Bridge?'

03 The MCP processes the request and returns structured, actionable information—whether that's a list of delays, a route map, or a sequence of steps.

The bottom line is you talk to your agent like talking to a human travel planner; it handles the complex data retrieval automatically.

Built For

Any role that relies on accurate, up-to-the-minute movement information. Think logistics coordinators running last-mile routes, software developers building smart city apps, or travel writers needing immediate status reports for clients.

Logistics Coordinator

Determines the best transport options for delivery personnel when road closures or subway delays make standard routing impossible.

Travel Blogger/Journalist

Quickly pulls live service status and journey plans to write real-time updates about London's public transit system without leaving the writing environment.

Software Developer

Integrates complex, multi-stage data points—like bus arrival prediction combined with road disruption alerts—into a functional prototype tool.

What Changes When You Connect

01 Avoid unexpected delays. Instead of waiting for a status board update, ask the agent to check the 'get_line_status' and know instantly if your commute is impacted.

02 Save time planning routes. The 'get_journey' tool handles complex transfers—it doesn't just give you directions; it gives you multiple options with cost estimates and changes counted.

03 Beat traffic bottlenecks. Need to drive across town? Use 'get_road_status' or 'get_road_disruptions' to know if the main artery is backed up before leaving your desk.

04 Manage transfers effortlessly. If a bus arrives late, use 'search_stop' first to find the correct ID, then 'get_arrivals' for precise predictions, linking real-world data to digital planning.

05 Integrate bike travel. Don't forget the last mile. Use 'get_bike_points' to check if there are available docks and bikes near your final destination.

Real-World Applications

A commuter needs alternative directions because of a subway issue.

The agent detects the Central line is suspended. The user asks for directions to Paddington. The agent uses 'get_journey' and calculates an optimal route combining overground trains with a necessary walk, bypassing the broken section entirely.

A tourist needs to know if their next bus connection is reliable.

The user asks for buses near Trafalgar Square. The agent uses 'search_stop' to get the correct stop ID, then calls 'get_arrivals'. It tells the tourist exactly when the 15 and 90 routes will arrive, preventing them from waiting unnecessarily.

A developer needs a reliable data source for an app prototype.

The dev uses the MCP to pull real-time road congestion using 'get_road_status' alongside tube line status via 'get_line_status', creating a single, comprehensive dashboard feed that updates automatically.

A bike messenger needs to check equipment availability.

Before starting a delivery run, the agent checks three nearby docking stations using 'get_bike_points'. It reports back: Station A has 5 bikes; Station B is full; and Station C has 12 available.

Patterns to Avoid

Trying to get status via a simple search.

✗ AVOID

Just searching 'London transport' in an AI client will give general articles about the Tube. It won't tell you if there are actual delays right now or what route is best for today.

✓ INSTEAD

You must use specific tools like 'get_line_status'. This tool connects directly to TfL's live feeds and gives a definitive answer on service disruptions, not just general information.

Asking for directions without specifying the start/end points.

✗ AVOID

Prompting: 'How do I get around London?' This is too vague. The agent can't plan anything because it doesn't know your origin or destination.

✓ INSTEAD

Always use 'get_journey' and provide specific locations (addresses, postcodes). For example: 'Plan a journey from Trafalgar Square to Canary Wharf.'

Assuming all bus stops are indexed by name.

✗ AVOID

Trying to get arrivals using just the general street name. The system needs precise location data; otherwise, it won't know which stop you mean and will fail.

✓ INSTEAD

First, run 'search_stop' using the area name. This gives you a valid Stop ID, which you then feed into 'get_arrivals' for accurate results.

The Right Fit

Use this MCP if your workflow requires real-time physical location data: knowing what's wrong with the Tube right now, checking road closures before sending out drivers, or confirming bus arrivals. It is built for actionable, time-sensitive movement intelligence.

Don't use it if you just need static information like Wikipedia facts about London or general historical travel guides. If your task involves scheduling appointments, using this MCP is overkill; a calendar tool does that better. Only rely on the tools when you absolutely need to know 'what is happening right now' in the physical transit network.

Keeping up with London's shifting commute schedule is exhausting.

Today, checking your journey involves hopping between three or four different websites: one for tube status, another for road traffic, and a third to see if the next bus is actually running. You spend time opening tabs, comparing colors (green vs. yellow), copying over stop IDs, and ultimately piecing together an answer that might still be slightly out of date.

With this MCP, you just ask your agent one question. It handles the entire sequence—checking line status, finding road disruptions, calculating transfers, and predicting arrivals—and hands you a single, cohesive set of directions. You get clarity, not tabs.

The Transport for London MCP provides reliable journey planning.

Manual planning requires confirming the route sequence using 'get_line_routes', then checking if that line is delayed via 'get_line_status'. If you miss a connection, you have to start over, manually cross-referencing bus arrivals with road status updates.

Now, your agent uses 'get_journey' and considers all these variables simultaneously. It tells you the best path immediately, whether that means taking the Piccadilly line or catching an alternate coach because of unexpected congestion.

Transport for London: 11 Tools

These tools let you access all facets of London's public transit data—from bus arrivals to major route planning—through a single, unified interface.

#	TOOL	DESCRIPTION
01	`get_arrivals`	Predicts live arrival times and details for buses at any specific stop point ID.
02	`get_bike_point_detail`	Retrieves detailed information about a single, specified bike docking station.
03	`get_bike_points`	Searches for nearby Santander Cycle stations and reports on available docks and bikes.
04	`get_journey`	Calculates travel plans between two points, offering multiple route options with costs and directions.
05	`get_line_detail`	Provides specific information about a selected Tube or Overground line (like the DLR or Elizabeth line).
06	`get_line_routes`	Lists the full sequence of stations served by a particular transit line.
07	`get_line_status`	Checks the current operational status of all major Underground lines, noting delays or suspensions.
08	`get_road_disruptions`	Lists details about ongoing road disruptions, including their cause and estimated time for clearance.
09	`get_road_status`	Shows the current congestion level (Good, Minor, or Severe) on major London roads.
10	`get_stop_details`	Retrieves detailed information for a specific bus stop or station ID.
11	`search_stop`	Finds matching London transport stops, whether they are bus stops or tube stations, using only the name.

See It in Action

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

U What's the status of the Central line?

Central line — Good Service. All trains running on time.

U Plan a journey from King's Cross to Heathrow.

Route 1: Piccadilly line direct (50min, £2.80, no changes). Route 2: Victoria to Green Park, then Jubilee to Bond Street, then Elizabeth line (45min, £3.40, 2 changes).

U When is the next bus at Oxford Circus?

Next buses at Oxford Circus: Route 8 to Holborn (3 min), Route 55 to Walthamstow (5 min), Route 98 to Willesden (7 min), Route 13 to Archway (9 min).

Frequently Asked Questions

01 **How do I check if a specific tube line is delayed using Transport for London MCP?**

You use the 'get_line_status' tool. Just ask your agent about the Central or Victoria lines, and it will report exactly what service status they have (Good Service, Minor Delays, etc.).

02 **Does Transport for London MCP help plan a trip involving multiple modes of transport?**

Yes. The 'get_journey' tool calculates multi-modal routes, combining tube travel with walking segments or even bus transfers into one step-by-step itinerary.

03 **Can I find out when the next bus is coming using Transport for London MCP?**

You can use 'search_stop' first to get a precise stop ID, and then feed that ID into 'get_arrivals'. This gives you predicted arrival times for specific routes.

04 **Is the road status accurate with Transport for London MCP?**

The 'get_road_status' tool checks major roads for congestion levels, while 'get_road_disruptions' provides details on unexpected closures and estimated times until they clear.

05 **What if I need to plan a bike-friendly journey in London?**

You can use the MCP in two steps: first, check availability with 'get_bike_points', then incorporate that starting point into your main route planning using 'get_journey'.

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 `"transport-for-london": { "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

Transport for London 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 →

INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by Transport for London. 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	Transport for London MCP
Server ID	`019d848c-2f4a-731f-aa55-771cb36bb030`
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/transport-for-london.