

M C P S E R V E R

N O C O D E

C L O U D H O S T E D

Consolidating Settlement Model for tailings.

Predict long-term settlement behavior and creep in mine waste.

Consolidating Settlement Model applies Terzaghi theory to model how mine tailings settle over time. Your AI client uses this MCP to calculate primary consolidation magnitude, predict settlement velocity, and account for long-term secondary compression caused by particle rearrangement.

tailings

consolidation

terzaghi

settlement

geotechnical-engineering



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.

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.

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.

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

Encryption at rest

Ed25519

PKI vault signatures

24h TTL

Ephemeral session keys

V8 Isolate

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.

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.

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.

06 — Honeypot Trap System

Phantom credentials are injected into isolated environments. If a honeypot 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.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

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.

Consolidating Settlement Model MCP

4 tools available

Cloud-hosted on Vinkius

You can use this MCP to handle the heavy math involved in geotechnical settlement analysis. Instead of manually running Terzaghi equations, you let your agent handle the calculations for primary and secondary consolidation. It's built specifically for mine tailings, where you need to account for both the initial rapid settlement and the much slower creep that happens as particles rearrange over time. You can get a clear picture of how a tailings deposit will behave from day one through years of long-term creep. This allows you to model settlement timelines and predict velocity without jumping between different calculation sheets. It's a direct way to bring geotechnical theory into your AI-driven engineering workflows.

Core Capabilities

01 — Primary Consolidation Modeling

Your agent uses this to find the initial magnitude and progress of settlement.

02 — Secondary Creep Calculation

The AI calculates long-term settlement caused by particle rearrangement after the primary phase ends.

03 — Settlement Velocity Prediction

Your agent determines the rate of settlement at specific time intervals.

04 — Timeline Generation

The AI builds a complete history of settlement over a requested period.

One Click on Vinkius — From Prompt to Execution

Available at vinkius.com/en/ai-agent-connect/consolidating-settlement-model — connect your AI agent in three steps.

- 01 Connect your AI client to the Vinkius hosted MCP.
- 02 Provide the tailings properties and time parameters to your agent.
- 03 The agent calls the relevant Terzaghi calculation tools.
- 04 The agent returns the settlement magnitude, rate, or timeline directly in your chat.

You connect the MCP to your client and start asking for specific settlement data.

Built For

Geotechnical engineers and mining professionals use this MCP to automate settlement calculations for tailings management.

Geotechnical Engineer

They hand off complex Terzaghi calculations to the agent to verify settlement models.

Mining Engineer

They use the settlement timelines to plan long-term tailings storage facility management.

Environmental Consultant

They use the creep data to assess the long-term stability of waste deposits.

What Changes When You Connect

- 01 Automates Terzaghi theory calculations within your AI client.
- 02 Separates primary consolidation from secondary creep for more accurate modeling.

- 03

Provides settlement velocity data at specific time steps.
- 04

Generates full settlement histories over custom time intervals.

Real-World Applications

Tailings Storage Design

Model how a new tailings deposit will settle to ensure facility stability.

Settlement Rate Monitoring

Determine the velocity of settlement to understand how quickly a site reaches stability.

Long-term Stability Assessment

Calculate secondary compression to predict how much a site will creep over years.

Consolidation Progress Tracking

Check the degree of primary consolidation at specific time intervals during site operation.

4 Tools Available

#	TOOL	DESCRIPTION
01	get_consolidation_rate	This tool calculates how fast settlement is happening at a specific point in time. It gives you the velocity of the settlement process.
02	get_primary_settlement	This tool handles the primary consolidation phase. It determines the total magnitude and how much progress has been made at any given time.
03	get_secondary_compression	This tool models the long-term creep that occurs after primary consolidation finishes. It calculates settlement caused by particle rearrangement.
04	get_settlement_timeline	This tool generates a high-level summary of total settlement. It maps out the settlement history across a series of time intervals.

See It in Action

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

U Calculate the primary settlement for a tailings layer with specific properties after 100 days.



The primary consolidation magnitude is 0.45 meters with a degree of consolidation of 85%.

U What is the total settlement including secondary compression after the primary phase is done?



The total cumulative settlement is 0.52 meters, consisting of 0.45 meters from primary consolidation and 0.07 meters from secondary compression.

U Generate a settlement timeline for the next 500 days.



The settlement timeline shows a rapid primary phase reaching 0.45m by day 150, followed by a slow secondary creep phase reaching 0.52m by day 500.

Frequently Asked Questions

01 What theory does this MCP use for calculations?

This MCP implements Terzaghi's consolidation theory to model both primary and secondary settlement.

02 Can I calculate long-term creep?

Yes, you can use the secondary compression tool to calculate settlement caused by particle rearrangement after primary consolidation is complete.

03 Which AI clients can I use with this MCP?

You can use this MCP with any compatible client like Claude, Cursor, or Windsurf.

04 Does this MCP handle both primary and secondary settlement?

Yes, it includes specific tools for primary consolidation magnitude and secondary compression creep.

05 How do I get a settlement history?

You can use the settlement timeline tool to generate a summary of total settlement over a series of time intervals.

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 <code>"consolidating-settlement-model": { "url": "..."} </code>
 Windsurf	MCP Settings → <code>mcp_settings.json</code> → 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

Consolidating Settlement Model 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 Consolidating Settlement Model. 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	Consolidating Settlement Model MCP
Server ID	01a08288-c412-7350-bc57-3361f84cf42a
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/consolidating-settlement-model.