

MCP SERVER

NO CODE

CLOUD HOSTED

# Subtitle Readability Pacer MCP for AI Agents

## Creating Accessible Video Captions That Prevent Viewer Fatigue

The Subtitle Readability Pacer calculates the optimal timing and block distribution for captions, ensuring viewers don't experience reading fatigue. It analyzes text density against human processing limits, enforcing strict thresholds like a 15 Characters Per Second (CPS) rate to maintain comfortable viewing accessibility.

**A+** Quality Score 100/100

subtitles

readability

video-processing

accessibility-tools

captioning



# 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](https://cloud.vinkius.com) — connect your AI agent in under 60 seconds.

# Subtitle Readability Pacer MCP

3 tools available

Cloud-hosted on Vinkius

When you work on video content, making sure subtitles are readable is half the battle. If captions move too quickly or contain blocks that are just too long, viewers get fatigued and tune out—no matter how good the original footage is. This MCP solves that problem by acting as a comprehensive readability check. It doesn't just give you character counts; it mathematically slices up your transcript into timed segments that match real human eye speed. You can run checks to gauge risk based on speaking rates, then use the system to automatically generate perfectly sized blocks for any duration. Vinkius hosts this MCP so your AI client connects once and gets immediate access to professional-grade captioning tools, saving you hours of manual timing adjustments.

---

## Core Capabilities

### 01 — Assess Caption Readability Risk

It assesses a given segment of text and tells you if it poses a reading risk based on established accessibility rules.

### 02 — Calculate Speaking Rate (CPS)

You can determine exactly how many characters per second are being spoken in any section of audio or text.

### 03 — Generate Timed Subtitle Blocks

It automatically splits long segments into multiple, timed subtitle blocks and assigns precise start and end times for each piece.

# One Click on Vinkius — From Prompt to Execution

Available at [vinkius.com/mcp/subtitle-readability-pacer](https://vinkius.com/mcp/subtitle-readability-pacer) — connect your AI agent in three steps.

- 01 You input the full text transcript and the total duration of the video segment you're analyzing.
- 02 The MCP analyzes this data, calculating the current speaking rate against professional standards (like the 15 CPS limit) and checking for block length violations.
- 03 It returns a detailed plan: either confirming low risk or providing multiple, precisely timed blocks that guarantee optimal readability.

The bottom line is you get a guaranteed, accessibility-compliant subtitle timeline without manual review.

---

## Built For

This MCP is essential for anyone working with video content who needs captions to be more than just text. If your job requires producing videos for large audiences—whether they're watching on mobile or in a professional setting—you need this tool. It prevents the frustrating experience of viewers getting lost because subtitles are too fast, saving you major headaches during final quality checks.

### Video Editor

You use it to process raw footage transcripts, ensuring that every caption block meets strict readability standards before the client sees the first cut.

### Content Accessibility Specialist

You run comprehensive checks on entire video libraries, proving compliance with accessibility guidelines by verifying ideal timing and character limits for all captions.

### Media Producer

You integrate it into your post-production workflow to quickly generate multiple versions of subtitles that fit different platforms' viewing requirements while maintaining viewer comfort.

## What Changes When You Connect

- 01 Stop worrying about viewer drop-off. By using the Subtitle Readability Pacer, you guarantee captions move at a comfortable pace, keeping viewers engaged through every minute.
- 02 Eliminate manual timing errors. The `generate_subtitle_blocks` tool automatically splits transcripts into timed chunks, saving hours of tedious review work on long videos.
- 03 Ensure compliance and comfort. You use the MCP to check for readability risks using `check_readability_risk`, guaranteeing your final product meets professional accessibility standards.
- 04 Accuracy over guesswork. Instead of guessing if text is too dense, you run `calculate_speaking_rate` to get a precise characters per second metric for every section.
- 05 Save time across entire libraries. Process multiple video files quickly by running batch checks that analyze density and duration simultaneously.

---

## Real-World Applications

### **A documentary needs captions, but the speaker talks too fast.**

The producer runs the transcript through the MCP's rate calculation tool. The system reports a high CPS risk level. They then use `generate_subtitle_blocks` to automatically split the rapid-fire dialogue into smaller, timed blocks that allow viewers time to read and absorb the information.

### **A corporate training video needs captions for global viewing.**

The content specialist uses the MCP to run a readability check on all sections. The system flags several blocks as potentially high risk due to long character counts. This allows them to revise only the problematic segments, ensuring compliance before final distribution.

### A podcast video needs captions for maximum accessibility.

The editor feeds the full transcript into the MCP. It assesses the entire run time and generates a perfectly timed sequence of blocks, guaranteeing that no line exceeds the 42-character limit or moves too quickly.

---

## Patterns to Avoid

---

### Just using word count as a proxy

#### X AVOID

Checking if a block has fewer than 30 words and assuming it's readable. This fails because the actual speaking rate (CPS) might still be too high, causing eye strain.

#### ✓ INSTEAD

Always use ``calculate_speaking_rate`` first. If that confirms an issue, use ``generate_subtitle_blocks`` to force optimal segmentation based on time and character limits.

---

### Manually timing every subtitle block

#### X AVOID

Spending hours in a video editor timeline, adjusting the start/end times for hundreds of caption boxes one by one. This is slow, prone to human error, and highly inefficient.

#### ✓ INSTEAD

Use ``generate_subtitle_blocks``. It takes your full text and duration and outputs perfectly timed, pre-formatted segments in minutes.

---

### Ignoring the character limit

#### X AVOID

Allowing blocks to run over 42 characters per line. This forces viewers to read huge walls of text that are difficult to track, especially on smaller screens.

#### ✓ INSTEAD

Run the ``check_readability_risk`` tool. It will flag these overly dense segments and suggest optimal splits using its other tools.

---

## The Right Fit

Use this MCP if your primary goal is accessibility compliance or maximizing viewer retention through readable captions. Specifically, if you need to enforce strict rules like the 15 CPS threshold or keep block lengths under 42 characters, this tool is non-negotiable. Don't use it if all you need is a simple word count or basic time stamping; those generic tools won't account for human processing speed. If your problem is simply 'I don't know what the video is about,' then

you just need transcription services—not this advanced timing engine.

---

---

## Subtitle Readability Pacer: Fixing Caption Timing in Video Accessibility

Today, many editors manually copy-paste transcripts into captioning software. They spend hours adjusting start and end times, constantly checking if the text is too long or if the pacing feels off to the naked eye. It's tedious work that often results in subtle readability flaws—a block that moves just a hair too fast, causing viewers to get frustrated before they even finish watching.

With this MCP, you feed it your full transcript and duration. It handles all the math: calculating speaking rates, checking density against professional standards, and then automatically generating perfectly timed subtitle blocks. You walk away with a fully compliant timeline that just needs approval.

---

## Subtitle Readability Pacer: Enhancing Viewing Comfort for Global Content

Manual review means you risk inconsistency across different video lengths or speaker styles. You might check one segment and miss a high-risk block in another part of the script, leaving your final product with accidental accessibility gaps.

This MCP runs systemic checks across the entire piece. It doesn't just fix blocks; it validates the whole narrative flow to ensure consistent readability from start to finish. Your content feels professional and polished every single time.

---

# Subtitle Readability Pacer: 3 Tools for Accessible Video Captioning

Use these tools to check for readability risks, calculate speaking rates, or automatically generate perfectly timed subtitle blocks for any video transcript.

#	TOOL	DESCRIPTION
01	<code>check_readability_risk</code>	Assesses if a specific subtitle text segment violates standard readability guidelines for viewers.
02	<code>calculate_speaking_rate</code>	Calculates the precise characters per second (CPS) rate of any given transcript section.
03	<code>generate_subtitle_blocks</code>	Slices a long text segment into multiple, timed subtitle blocks that adhere to character limits and time duration.

---

## See It in Action

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

**U** Hey, can you check if this 45-second segment of dialogue is safe for captions?



### Caption Readability Assessment

- **Total Characters:** 720
- **Duration:** 45 seconds
- **Calculated Speaking Rate (CPS):** 16.0 CPS
- **Risk Level:** HIGH RISK

> The rate exceeds the optimal 15 CPS threshold, leading to potential viewer fatigue. You need segment splitting.

**U** I have a transcript of 180 characters that lasts exactly 6 seconds. What should I do?



### Analysis Complete

- **Rate:** 30 CPS (Too fast)
- **Recommendation:** Split this into two blocks.

Block	Start Time	End Time	Characters
1	0.0s	2.5s	~60 chars
2	2.5s	6.0s	~60 chars

*This keeps the reading rate below the recommended limit.*

- U** Generate subtitle blocks for this chunk that is 150 characters and runs for 9 seconds.



### Subtitle Blocking Plan

The segment has been optimally split into three parts:

- **Block 1:** (Start: 0ms, End: 3000ms) - ~45 chars
- **Block 2:** (Start: 3000ms, End: 6000ms) - ~50 chars
- **Block 3:** (Start: 6000ms, End: 9000ms) - ~55 chars

This timing keeps all blocks under the 42-character limit and maintains optimal viewer pacing.

---

## Frequently Asked Questions

---

### 01 How does Subtitle Readability Pacer help me with captioning compliance?

It ensures your captions meet professional accessibility standards. The MCP checks for ideal timing, confirming that the speaking rate and block lengths are within comfortable limits, so viewers don't get fatigued or confused.

### 02 I have a long video; can Subtitle Readability Pacer handle all the blocks?

Yes. You feed it the full transcript and duration. The system analyzes the entire piece and automatically generates hundreds of perfectly timed subtitle blocks, saving you massive amounts of manual editing time.

### 03 What if my speaker talks fast? Will Subtitle Readability Pacer fix it?

It detects overly fast speaking rates. If the rate is too high, the MCP recommends or generates multiple smaller blocks to slow down the visual pacing and make the information digestible for viewers.

### 04 Is this better than just manually timing my subtitles?

Absolutely. Manual timing is prone to human error and inconsistency. This MCP uses mathematical models to guarantee that every single block adheres to best practices, ensuring consistent quality across your entire video library.

### 05 Do I need this for corporate training videos specifically?







Yes, especially. Training content needs to be crystal clear and accessible. This MCP validates that the pace of the captions supports learning retention by optimizing timing and block size.

# Go Live in 60 Seconds

Get your connection token from [cloud.vinkius.com](https://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
 <b>Claude AI</b>	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
 <b>Cursor</b>	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
 <b>VS Code</b>	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add <code>"subtitle-readability-pacer": { "url": "..." }</code>
 <b>Windsurf</b>	MCP Settings → <code>mcp_settings.json</code> → Add endpoint URL
 <b>ChatGPT</b>	Settings → Tools & plugins → Add MCP server → Paste endpoint
 <b>Gemini</b>	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

# Subtitle Readability Pacer 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](https://cloud.vinkius.com) →

[vinkius.com](https://vinkius.com) · [support@vinkius.com](mailto: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 Subtitle Readability Pacer. 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	July 2026
MCP Server	Subtitle Readability Pacer MCP
Server ID	019f2510-e0d1-71bd-acea-15edfbb85931
Platform	Vinkius Cloud for AI Agents
Endpoint	<a href="https://edge.vinkius.com/{token}/mcp">https://edge.vinkius.com/{token}/mcp</a>

### 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/subtitle-readability-pacer](https://vinkius.com/mcp/subtitle-readability-pacer).