MCP SERVER | NO CODE | CLOUD HOSTED

Presentation Duration Calculator MCP

Get accurate timing and pacing for every slide.

Presentation Duration Calculator MCP helps you plan exactly how long your talk will last. Instead of guessing, you can use your AI client to calculate total run times based on slide counts or get specific pacing advice for different presentation styles like technical deep dives or quick pitches.

A+ Quality Score 100/100

presentation | timing | planning | productivity | scheduling

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · SEPTEMBER 2026
vinkius.com/en/ai-agent-connect/presentation-duration-calculator

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.

Presentation Duration Calculator MCP

4 tools available
Cloud-hosted on Vinkius

Planning a presentation usually involves a lot of guesswork regarding how much time you actually need. This MCP removes that uncertainty by giving your AI client the math it needs to build a realistic schedule. You can determine your total runtime by providing a slide count and a preferred pace, or you can ask for specific pacing recommendations based on the type of talk you are giving. Whether you are prepping a high-speed pitch or a dense technical walkthrough, the tool provides the logic to keep you on track. It also lets you see a high-level summary of your planned session so you can spot timing issues before you step on stage. It is a practical way to ensure your content fits your allotted slot without rushing through critical points or dragging on too long.

Core Capabilities

01 — Runtime Estimation

Your agent calculates total duration based on slide volume.

02 — Style-Based Pacing

The MCP suggests how much time to spend per slide for specific talk types.

03 — Session Summarization

Your AI client generates a summary of the presentation's overall characteristics.

04 — Tier Discovery

The tool lists all available presentation tiers to help you choose a style.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/en/ai-agent-connect/presentation-duration-calculator** — connect your AI agent in three steps.

01 Connect the MCP to your preferred client like Claude or Cursor via Vinkius.

02 Tell your AI client how many slides you have and your intended pace.

03 Ask for a pacing recommendation if you aren't sure how fast to go.

04 Review the calculated duration or summary to finalize your plan.

Connect the MCP to your client and start asking for timing data.

Built For

This is for anyone who needs to manage time during a live talk or meeting. It helps professionals move from rough ideas to a timed agenda.

Public Speakers

They hand off slide counts to the AI to ensure they don't run over their time limit.

Project Managers

They use the tool to schedule meeting blocks that match the content being presented.

Technical Trainers

They request pacing recommendations for deep dive sessions to keep students engaged.

What Changes When You Connect

01 Eliminates guesswork in presentation scheduling.

02 Matches pacing to the specific intent of the talk.

03 Provides a quick overview of session characteristics.

04 Prevents running over allotted time slots.

Real-World Applications

Pitch Deck Preparation

Calculate exactly how much time you have for each slide during a high-stakes investor pitch.

Technical Workshops

Get pacing advice for deep dive sessions to ensure complex topics get enough time.

Meeting Planning

Determine if your planned presentation fits within a standard 30 or 60 minute meeting window.

Conference Scheduling

Summarize presentation characteristics to help organizers manage event timelines.

Patterns to Avoid

Guessing your total runtime

AVOID

You tell your agent you have 20 minutes and hope your 30 slides fit. You end up rushing the conclusion because you didn't do the math.

✓ INSTEAD

Provide your slide count to `get_estimated_duration` to get a realistic time estimate before you start building slides.

Using a single pace for everything

AVOID

You apply the same speed to a quick elevator pitch and a deep technical workshop. One feels rushed while the other drags.

✓ INSTEAD

Use `get_pacing_recommendation` to find the right speed for your specific talk type.

Ignoring available pacing styles

AVOID

You pick a random speed for your presentation without checking if there is a standard category for your specific goal.

✓ INSTEAD

Run `list_presentation_tiers` first to see which pacing categories match your intended presentation style.

The Right Fit

Connect this MCP if you regularly manage live talks, workshops, or pitches where timing is critical. Don't connect it if you are just drafting casual notes or unstructured content. The decision depends on whether you have a fixed time slot to fill.

Presentation Duration Calculator Tools

These tools handle the math and pacing logic for your presentation planning. They turn slide counts into actionable schedules.

#	TOOL	DESCRIPTION
01	`get_estimated_duration`	Calculates the total length of your presentation. You provide the slide count and pacing to get a final time estimate.
02	`get_pacing_recommendation`	Provides ideal time-per-slide suggestions. It suggests speeds based on specific presentation tiers like a Quick Pitch or Technical Deep Dive.
03	`get_presentation_summary`	Generates a high-level overview of your session. It summarizes the characteristics of your planned presentation.
04	`list_presentation_tiers`	Shows all available presentation styles. Use this to see which pacing categories are available for your planning.

See It in Action

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

U How long will a 20-slide presentation take if I spend 3 minutes per slide?

The total duration for a 20-slide presentation with 3 minutes per slide is 60 minutes.

U What is the recommended pacing for a Technical Deep Dive?

For a Technical Deep Dive, the recommended pacing is 10 minutes per slide.

U Give me a summary for 10 slides at 5 minutes each.

The presentation will last 50 minutes, which is considered a Long Form presentation.

Frequently Asked Questions

01 **How do I use this MCP with my AI client?**

You connect to the MCP through the Vinkius platform. Once connected, your AI client can immediately use the tools to calculate timing and pacing.

02 **Can I get different pacing for different types of talks?**

Yes. You can use the pacing recommendation tool to get specific advice for different tiers, such as a Quick Pitch or a Technical Deep Dive.

03 **Does this tool calculate the total time for me?**

Yes. The get_estimated_duration tool calculates the total time based on your slide count and chosen pace.

04 **What kind of summary can the AI provide?**

The tool provides a summary of the presentation characteristics, including the total duration and the presentation type.

05 **Do I need to provide the slide count?**

Yes, you provide the slide count to the AI client so it can use the tool to calculate the total duration.

06 **How do I know how much time to spend on each slide?**

You can use the `get_pacing_recommendation` tool to receive a suggested time per slide based on your presentation tier, such as an Educational Session or a Standard Briefing.

07 **Can I include extra time for Q&A?**

Yes, when using `get_estimated_duration`, you can provide a `bufferMinutes` value to account for questions, technical delays, or transitions.

08 **What is a 'Long Form' presentation?**

A presentation is categorized as Long Form by the `get_presentation_summary` tool if the total calculated duration is 45 minutes or more.

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 `"presentation-duration-calculator": { "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

Presentation Duration Calculator 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 Presentation Duration Calculator. 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	Presentation Duration Calculator MCP
Server ID	`01a0cbef-3d75-729e-b9d9-d24d7a8fd854`
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/presentation-duration-calculator.