



# DECKRUN

A document-to-deck engine for AI agents and backend services, delivered as a paid subscription service on a cloud platform. It turns structured markdown into polished PDFs, narrated videos, and audio assets, with enterprise-grade connectors and async job handling – all operated for you.



## WHY DECKRUN

AI systems can generate high-quality content — presentations, documents, videos, datasheets — but **they cannot finish the work**. Today, teams copy AI-generated content into slide tools, video editors, and document software to make it usable. The moment they do, two things happen:

- 1. Execution slows down:** Work moves across tools, licenses, and manual steps just to turn content into deliverables.
- 2. AI memory breaks:** Once content is exported, it becomes opaque to the AI that created it. Over time, LLMs can no longer reliably answer:
  - What was delivered to this customer?
  - What issues were already addressed?
  - How has messaging evolved over months or quarters?

Deckrun solves both problems.

Deckrun is the **execution layer for AI-generated content**. It turns LLM output directly into finished artifacts — PDFs, decks, videos, and narrated assets — **without leaving the AI system**.

Because outputs are generated from structured schemas and preserved with metadata and provenance:

- Content remains **machine-readable and queryable**
- AI systems can reason over historical deliverables
- Knowledge compounds instead of fragmenting

The result:

- 80%+ of content use cases delivered instantly
- Fewer tools, licenses, and manual steps
- A persistent, AI-consumable record of what was created, shared, and said

Deckrun doesn't just make AI faster — it allows AI to **remember**.

## PRODUCT OVERVIEW

Deckrun is **agent-native and schema-driven**: the same markdown and theme produce the same outputs every time, so decks are reproducible and safe to call at scale from LLM-based systems.

ATTRIBUTE	DESCRIPTION
<b>Primary inputs</b>	Markdown + theme + assets (images, logos, etc.)
<b>Primary outputs</b>	PDF slide decks, MP4 video (slides + timing, optional TTS), audio narration (MP3), notes text, and zip bundles for connectors.

### Use Cases

- ✓ Webinars with cloned narrator voice
- ✓ AI docs → PDF, video, audio from one source
- ✓ Resumes & multi-layout from one markdown
- ✓ Marketing & sales collateral
- ✓ YouTube & tutorial videos
- ✓ Podcasts (audio-only; higher tiers)

### Demos

Watch Deckrun in action on [YouTube: @AgenticDecks](#):

- **Branded decks** — Logo, theme, footer; Deckrun default and custom themes
- **Narration** — System voice, voice cloning, slide notes → audio/video
- **Layouts** — Title, two-column, section headers, images, timers; full layout showcase

ATTRIBUTE	DESCRIPTION
<b>Consumption model</b>	HTTP API and backend workflows (e.g. chatbots, automation).
<b>Commercial model</b>	Paid subscription only (no free tier), with usage and feature tiers for different customer profiles.

## KEY BENEFITS

- **Agent-centric:** Request/response schemas, zip bundling, and async jobs are explicitly designed for LLM agents and orchestration frameworks.
- **Deterministic and inspectable:** Markdown + theme + schema; every deck can be reproduced and audited.
- **Reduced operational overhead:** One markdown source → PDF, video, notes, and audio; repeatable webinar and launch content without manual slide or video tooling.
- **Enterprise-ready:** Runs in your environment; integrates with enterprise storage and messaging via connectors.
- **Future-proof:** Audio, premium voice cloning, and user-provided narration support keep visual and audio experiences synchronized over time.

## CORE CAPABILITIES

- **High-fidelity PDF generation**
  - Markdown → PDF pipeline with themeable layouts.
  - Slide templates (title, two-column, image-heavy, etc.) and brand-safe colors/typography; one source → resumes (single- or two-column), pitch decks (internal vs external, detailed vs summary).
  - Support for embedded images and assets (base64 or uploaded files). Themes support branding and personalization: logo, accent colors, imagery, footer/contact, and persona per customer.
- **Video output (slides → MP4)**
  - Turn any markdown-backed deck into a ready-to-publish video; ideal for webinars, YouTube tutorials, and launch content.
  - Automatically paces slides using timers or sensible defaults.
  - Video is rendered from the same deterministic source as the PDF, eliminating visual drift and narration mismatches.
  - Every run gives you links to the video, PDF, and notes for reuse across channels (marketing collateral, sell sheets, walkthroughs).
  - **Embedded MP4 clips:** Drop a recorded demo, product walkthrough, or screen capture directly into a slide using `video_src`. The clip is stitched into the output MP4 at that position — trimmed to a window if needed — while surrounding slides remain narrated or timed as usual. In PDF mode the slide renders normally; the clip tag is ignored.
- **Notes text extraction**
  - Every slide's presenter notes are captured as clean, reusable text for TTS (including founder or system voice), blog posts, summaries, FAQs, or agent analysis. AI can generate slide content and notes; Deckrun renders PDF, video, and audio synced to slides.
- **Async job model for heavy outputs**
  - Large videos and audio tracks are generated in the background; your system receives a job ID and checks back for status and asset links.
  - Safe to call from chatbots, workflows, and integrations without timeouts. All heavy outputs follow this async job model.
- **Output connectors + zip bundling**
  - Deliver finished assets to cloud storage (S3, Google Drive, Dropbox, Box), internal filesystems, or Slack.
  - Optionally bundle into a single zip (PDF, video, notes, audio). Connectors and bundling are handled by the service. Agents (or users) generate markdown; Deckrun produces PDFs, decks, video, and optional audio from the same version-controlled source.
- **Agent-friendly schemas**
  - Strict, well-documented schemas for slides, themes, job status, and outputs so platform and AI teams can validate requests and integrate Deckrun as a predictable rendering step.

## ARCHITECTURE & DEPLOYMENT

- **Core engine**

- API-first service built for agent workloads.
- High-performance rendering pipeline for consistent PDF output.
- Built-in video composition and an async job engine for heavy workloads (video, audio).
- **Audio pipeline**
  - Slide notes → TTS (system voices or premium AI voice cloning) → per-slide audio segments → a single polished MP3. One markdown script → branded PDF deck, narrated video (founder or system voice), speaker notes and audio for campaigns.
  - Timing map so audio and visuals stay in lock-step; chapter markers for supported players.
  - Supports: system TTS, high-performance TTS for voice cloning from a reference sample, and user-provided audio + CUE (mux custom narration with generated video).
  - Designed for founder- or executive-voice narration across webinars, product launches, and recorded content. Podcast-style audio-only versions of presentations are available in higher tiers.
- **Deployment:** Cloud platform, operated for you; enterprise options for deeper integration and tighter controls.

## INTEGRATION & INTERFACES

---

- **HTTP API:** Primary interface for submissions, job status, and asset retrieval.
- **Schemas:** Well-documented request/response schemas for slides, themes, job status, and outputs; validation before production. Typical flows: agent-generated markdown → PDF + video + notes; single source for resumes, pitch decks, and webinar content.
- **Connectors:** Deliver outputs to S3, Google Drive, Dropbox, Box, internal filesystems, Slack; expanded coverage and custom connectors in higher tiers.
- **Output formats:** PDF, MP4, MP3, notes text, zip bundles (PDF + video + notes + audio).

## SECURITY & COMPLIANCE

---

- **Authentication / access:** API keys for service access; SSO/SAML integration available on Enterprise plan.
- **Data handling:** Processing and storage in accordance with subscription terms; enterprise plans support custom security reviews and coordinated rollout.
- **Compliance:** Enterprise plans include contractual SLAs and security review options; contact for specific compliance posture (e.g. SOC2).

## PERFORMANCE & SCALABILITY

---

- **Throughput:** Tiered by plan; Growth for bounded concurrency and async job queue, Pro for higher throughput and larger concurrent job limits, Enterprise for platform-scale needs.
- **Async by default:** All heavy outputs (video, audio) use the async job model; no long-lived request timeouts.
- **Concurrency:** Bounded concurrency and job queue in Growth; higher limits in Pro and Enterprise.

## COMMERCIAL MODEL

---

Deckrun is offered exclusively as a **paid subscription**. Plans are tiered by monthly output volume and capability access; advanced outputs are available as add-ons so base plan prices stay nominal.

### SUBSCRIPTION PLANS

- **Starter** For developers evaluating Deckrun or small teams with occasional PDF output needs. Includes: PDF generation and notes text extraction.
- **Starter Plus** For small teams running one or two agent workflows with regular content output. Adds: System TTS narration and silent video (slides → MP4).
- **Growth** For teams embedding Deckrun across multiple agent workflows. Adds: Output connectors (S3, Google Drive, Slack, filesystem) and the async job queue for heavy workloads.
- **Pro** For organizations standardizing Deckrun across multiple products or business units. Adds: Higher throughput, priority job queue, and expanded connector coverage.
- **Enterprise** For platform teams and regulated environments. Custom engagement — connectors, security reviews, SSO/SAML, and contractual SLAs.

### ADD-ONS (AVAILABLE ON ANY PLAN)

- **Premium Voice narration** — High-fidelity expressive or cloned voices powered by a professional TTS provider, for customer-facing videos, launches, and branded assets. Available in Standard and Pro capacity tiers.

- **Video Splice** — Embed recorded MP4 clips (product demos, screen captures, b-roll) directly into a slide position in the output video. The clip is transcoded and stitched into the timeline at that point; the rest of the deck continues as narrated or timed slides. Available in Standard and Pro capacity tiers.

All plans and add-ons include access to the same **APIs and schemas**. Plans govern output volume and base capabilities; add-ons unlock specialized output types used by a subset of customers.

#### USAGE-BASED PRICING WITH RENDER UNITS

**Deckrun pricing is based on what gets rendered — not on seats or file types.** Each job consumes Render Units based on the number of pages and the outputs produced (PDF, video, narration, and machine-readable metadata). This allows teams to generate multiple high-quality artifacts from the same source content **without paying repeatedly for the same work.**

#### WHAT THIS PRODUCT IS NOT

- Not a slide editor or visual design tool.
- Not a consumer presentation app.
- Not an LLM writing system — Deckrun **renders and orchestrates** outputs from structured inputs your agents or tools generate.

#### GETTING STARTED / NEXT STEPS

Deckrun is currently offered via direct engagement. Typical onboarding includes:

- A short technical walkthrough of your use case (e.g. webinar generation with founder voice)
- Validation of output quality and throughput requirements
- Trial access under a paid subscription tier

If this aligns with your workflow, we can move quickly to a pilot.

[LEARN MORE](#)

To learn more, visit [agenticdecks.com](https://agenticdecks.com)

© 2026 Agentic Decks. All rights reserved.

Built with Deckrun

