A TECHNICAL AND COMMERCIAL DOSSIER

Autonomous Video Studios

The original faceless-channel programme and the ChannelForge SaaS: design, operations, results and economics



Frames from six of the most-watched videos, as they were posted

2,959
videos made by the two systems

4.43M
views on the original's channels

65 + 49
channels: original + ChannelForge

Majid Khan Mohmand

linkedin.com/in/majidkhanmohmand · fiverr.com/majidkhan_moh · WhatsApp +39 331 106 0553

Via Giuseppe Mazzini 4, 95032 Belpasso (CT), Italia

Read it online: dossier.storeswiki.com

Second edition · 14 November 2026

A TECHNICAL AND COMMERCIAL DOSSIER

Autonomous Video Studios

The original faceless-channel programme and the ChannelForge SaaS: design, operations, results and economics

Written by

Majid Khan Mohmand

LinkedIn	linkedin.com/in/majidkhanmohmand
Fiverr	fiverr.com/majidkhan_moh
WhatsApp	+39 331 106 0553
Address	Via Giuseppe Mazzini 4, 95032 Belpasso (CT), Italia
Live systems	content.storeswiki.com · forge.storeswiki.com

READ IT ONLINE

dossier.storeswiki.com

The same dossier as a website: every chapter, table, chart and diagram, searchable, for any screen.

Second edition · 14 November 2026. Figures are measured from the live systems on 14 November 2026 at 17:09 UTC (18:09 in Rome) unless marked as estimates. The first edition was written on 3 November 2026; everything the systems gained since is in this one.

About this document

This dossier describes two production systems that design, make, check, publish and measure short-form and long-form videos without a person in the loop: the **original programme**, which runs 65 channels of its own owner, and **ChannelForge**, the same engine rebuilt as a multi-tenant software-as-a-service product with workspaces, teams, credits, billing and an API. It is written for someone who may want to acquire either system, licence it, or have one built for them, and it tries to answer every question such a person would ask: what the systems do, how they are built, what they have produced, what they cost to run, what you would receive and what you would need.

Every number in it comes from the systems' own records — their databases, ledgers and logs — read on 14 November 2026. Where a figure is an estimate (monthly running costs, for example) it is labelled as one. Audience figures are what the platforms reported to the systems' analytics at that time; they grow every day.

Screenshots of the two web applications come from their development builds, which run on demonstration data so that no account or customer information is shown; the published pages of ChannelForge are captured live. Internal addresses, account identifiers and credentials are deliberately left out: they are part of the handover, not of a document that may be forwarded.

How to read it

- **Part I** is the overview: what the systems do and the headline numbers. Read it in ten minutes. The page after this one lists what is new since the first edition.
- **Part II** explains the original programme: its channels, how an episode is made, the Blender episodes, captions that are the voice, films, publishing, every failure and its lever, self-healing and the portal.
- **Part III** covers ChannelForge: the product, what it gained since the first edition, its architecture, tenancy and security, GPU allocation, plans and the API.
- **Part IV** is infrastructure, operations and economics: the cloud layout, the GPU fleet, backups and disaster recovery, the technology stack, costs and the build timeline.
- **Part V** shows the results and sets out what a buyer receives, followed by the author's other products and three appendices.

© 2026 Majid Khan Mohmand. All rights reserved. The systems, their code, channels, brands and media described here are the author's property. This document may be shared with prospective buyers and partners in full and unaltered.

THE SECOND EDITION

What changed in eleven days

The first edition described both systems as they stood on 3 November 2026. In the eleven days since, they gained more than in any stretch before — 1,146 commits — and their output more than doubled. Every figure in this edition was read again on 14 November 2026 at 17:09 UTC (18:09 in Rome).

	3 November	14 November	
Episodes made by the original programme	941	2,015	×2.1
Posts by the original programme	2,261	4,421	×2.0
Views on its live channels	2.15 million	4.43 million	×2.1
Followers and subscribers	5,279	13,481	×2.6
Channels (live)	33 (13)	65 (16)	
Long-form films	16	28	
ChannelForge episodes	440	944	×2.1
ChannelForge channels	22	49	
ChannelForge features switched on	36	166	×4.6
GPU boxes (original)	12	30	
Lines of code · test files	700,000 · 450	1.14 million · 1,106	

What is new

- **Blender episodes.** Photoreal 3D scenes built from free CC0 assets: two of every live channel's daily slots in the original, a choice of picture on ChannelForge. *Chapter 'Blender episodes'*.
- **Captions that are the voice.** Every spoken word on screen, timed by a speech recogniser that hears the finished narration, in every project's shorts. *Chapter 'Captions that are the voice'*.
- **32 new channels in four new projects**, and three new ChannelForge projects in the owner's workspace; three more channels live. *'Ten projects', Appendices A and B.*
- **Every failure has a lever.** A ledger of every failed post, and beside each the button that fixes it, in both systems. *Chapter 'Every failure has a lever'*.
- **Every page in tabs.** Both web applications rebuilt page by page, nothing lost; a Connections page and a live schedule. *'The content portal', 'What ChannelForge gained'*.
- **162 new ChannelForge features**, all but two from three product studies; 130 are on in the owner's workspaces, proven there first. *Chapter 'What ChannelForge gained since the first edition'*.
- **Stronger models, a bigger fleet.** Claude Fable 5.2 first in a five-model chain; 30 GPU boxes in 13 pools that fail over across zones, regions and cloud accounts. *'How an episode is made', 'The GPU fleet'*.
- **Films, held and rescued.** A fairer release floor, picture fixes by hand, and films for channels not yet live. *Chapter 'Long-form films'*.
- **Everything twice.** Every video copied hourly to a second bucket, the code to GitHub every four hours. *Chapter 'Media, backups and disaster recovery'*.

AT A GLANCE

Six reasons to look closer

2,959

videos in 33 days

made, checked and published with no one filming, editing or uploading: 2,015 by the original programme, 944 on ChannelForge.

4.43M

views, without paid promotion

on 16 live channels in their first month, with 13,481 followers and subscribers — ×2.1 in the eleven days since the first edition.

8.08

average critic score, out of 10

an episode under 8 never reaches a platform: 30 of 2,015 were held by the gates, and a fact, a unit or a caption out of sync holds one too.

$7.02

a short on ChannelForge, in model time

charged to customers at cost times 1.8 to 3 — a margin of 44 to 67% on every credit at list value, so every episode a customer makes earns more than it costs.

490 of 521

failed posts recovered

every refusal on the record with a lever beside it; GPU work fails over across zones, regions and cloud accounts, as it did on 14 November.

1.14M

lines of code, handed over turn-key

1,106 test files, golden GPU images, runbooks and a written disaster-recovery plan; either system can be sold on its own.

33 days, from the first episode

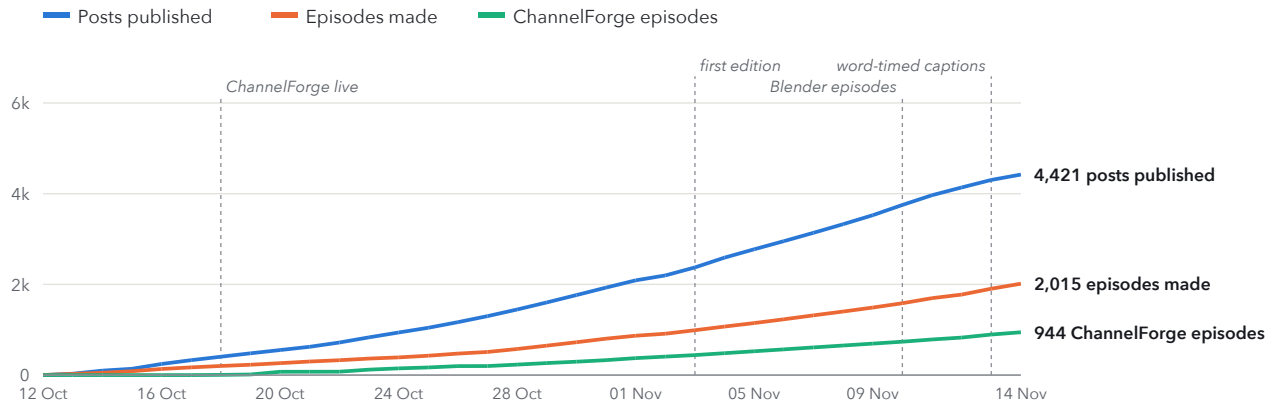


Figure 1. Everything the two systems made and published, added up day by day (12 October – 14 November 2026): the original programme's posts and episodes and ChannelForge's episodes, with the days that changed the pace.

EXECUTIVE SUMMARY

Two video studios that run themselves

In 33 days, from 12 October to 14 November 2026, the two systems described here made **2,959 finished videos** and published them as **4,621 posts**. The original programme's channels have been watched **4,434,425 times** — more than twice the 2.15 million of the first edition, eleven days earlier. Nobody films, edits, writes or uploads anything: an AI agent writes each episode — as code, or since 10 November as a photoreal Blender scene — a GPU renders it, a voice model narrates it, captions are timed to every spoken word, a battery of quality gates checks it, and the system posts it on a human-looking schedule to each channel's own accounts — then reads the audience numbers back and learns from them.

2,015
episodes made by the original programme

4,421
posts on YouTube, Facebook, Instagram, LinkedIn, TikTok

4.43M
views across the original's 16 live channels

28
long-form films made

944
episodes made on ChannelForge

2,735
ChannelForge jobs run

166
ChannelForge features switched on, of 198

117
Blender episodes since 10 November

The original programme

Ten projects and 65 channels in science, engineering, nature, sport, craft, history and the human body. 16 channels are live on YouTube, Facebook and Instagram, with a curated LinkedIn queue and TikTok (X is connected and waits for API credits); the other 49 generate every day towards launch. Each live channel posts two to four times a day — two of its slots now Blender episodes — a long-form film goes out every day, and a self-healing layer watches every slot, every GPU and every platform, making up anything that was missed and putting a button beside every failure it cannot fix alone.

ChannelForge

The same engine as a product: anyone can describe a niche, receive channel proposals, have a brand and a renderer built, and let an autopilot make and publish episodes to nine platforms, with Mastodon and Telegram built as the tenth and eleventh. It is multi-tenant (workspaces, projects, teams, roles, approvals), metered (credits charged at cost times a plan markup), billed (Stripe and PayPal built in, switched on at launch), and open (REST API, signed webhooks, an MCP server for AI agents). 166 of its 198 features are switched on — 36 for every customer, from a hook lab and language twins to retention learning, and the rest, most of them from three studies of leading products, in the owner's own workspaces first. It makes Blender episodes too, and keeps a ledger of every failed post with a lever beside each.

How it is built

Python control planes, TypeScript/Next.js web applications, PostgreSQL, Redis, headless Chrome with WebGL2 and Blender 4.5 for rendering, ffmpeg, Claude Fable 5.2 and 5.1, Opus 5.5 and Sonnet 5.5 through Amazon Bedrock, Google Chirp 3 HD voices, a speech recogniser on every GPU that times each spoken word, NVIDIA L4 GPUs on AWS in two accounts and two regions, Google Cloud for hosting and the media archive, Cloudflare at the edge. Everything is in Git, backed up to GitHub every four hours with a secret scanner in front, and recoverable from a written turn-key runbook.

Economics

A short episode costs about **$10.82** in model time on the original programme on average — about $10 while its agent was Claude Opus 5.5, about $16 since it moved to the stronger Claude Fable 5.2 on 10 November — and **$7.02** on ChannelForge, whose standard agent is Claude Sonnet 5.5, plus one to two dollars of GPU time. ChannelForge sells credits from $49 a month and charges every job at its cost times a plan markup of 1.8 to 3, so every episode a customer makes earns more than it costs.

> **WHAT IS ON OFFER**
>
> Either system – or both – can be acquired outright, licensed, or delivered as a done-for-you build on the buyer's own cloud. The handover includes all source code, infrastructure scripts and golden machine images, the complete documentation and runbooks, the deployed-state snapshot, and a guided transfer. The chapter 'What a buyer receives' lists everything; the one after it introduces the author.

Key facts

Built	12 October – 14 November 2026: 33 days from the first episode to this edition
Original programme	10 projects, 65 channels; 16 live on YouTube, Facebook and Instagram, with LinkedIn and TikTok
ChannelForge	a multi-tenant SaaS: 9 platforms open and 2 more built, 166 of 198 features on, plans from $49 a month, a REST API and an MCP server
Made so far	2,015 (117 of them in Blender) + 944 short episodes and 28 + 3 long-form films (original + ChannelForge)
Published	4,421 posts by the original programme, 200 by ChannelForge's workspaces
Audience	4,434,425 views and 13,481 followers on the original's live channels
AI models	Claude Fable 5.2 and 5.1, Opus 5.5, Opus 5, Opus 4.8, Sonnet 5.5 on Amazon Bedrock; Google Chirp 3 HD voices
GPUs	30 NVIDIA L4 boxes for the original in 13 pools, plus ChannelForge's own; two AWS accounts, two regions, on only while there is work
Hosting	Google Cloud (Belgium): two containers, an Arm VM, Cloud SQL, 730 GB of media in Cloud Storage, copied hourly to a second bucket
Code	about 1.14 million lines in three trees, 1,106 test files, 1,527 commits, backed up to GitHub every four hours
Recovery	a turn-key runbook (DR.md), a nightly snapshot of everything deployed, a posting guard
Status	both systems live; ChannelForge's selling switched off until its launch

Contents

New in the second edition . . . iv
At a glance . . . v
Executive summary . . . vi

PART I · THE SYSTEMS AT A GLANCE 9
1 What the two systems do . . . 10
2 The numbers so far . . . 13

PART II · THE ORIGINAL PROGRAMME 15
3 Ten projects, 65 channels . . . 16
4 How an episode is made . . . 19
5 Blender episodes: a second kind of picture . . . 21
6 Captions that are the voice . . . 24
7 Long-form films . . . 26
8 Publishing and the calendar . . . 28
9 Every failure has a lever . . . 31
10 Self-healing operations . . . 33
11 The content portal . . . 35

PART III · CHANNELFORGE 38
12 The product . . . 39
13 What ChannelForge gained since the first edition . . . 43
14 Architecture . . . 47
15 Multi-tenancy, security and isolation . . . 49
16 GPU allocation . . . 51
17 Plans, credits and billing . . . 52
18 API, webhooks and agents . . . 54

PART IV · INFRASTRUCTURE, OPERATIONS AND ECONOMICS 55
19 Infrastructure map . . . 56
20 The GPU fleet . . . 58
21 Media, backups and disaster recovery . . . 60
22 Technology stack . . . 62
23 Costs and unit economics . . . 63
24 Timeline and effort . . . 65

PART V · RESULTS AND THE OFFER 66
25 Audience and results . . . 67
26 What a buyer receives . . . 72
27 About the author and the ecosystem . . . 73

APPENDICES 75
A Channel directory . . . 76
B ChannelForge's workspaces and channels . . . 79
C Glossary . . . 81

PART I

The systems at a glance

What the two systems do, who they serve, and what they have produced in their first month.

CHAPTER 1 · OVERVIEW

What the two systems do

Faceless channels — channels whose videos carry a voice and animation but no presenter — are one of the fastest-growing formats on YouTube Shorts, Instagram Reels and Facebook. They are also expensive to run by hand: a good one needs research, a script, animation, narration, editing, captions, thumbnails, posting at the right hours on several platforms, replies to comments and a constant eye on what works. These systems do all of it, continuously, for dozens of channels at once.

An episode, in one paragraph

A topic planner proposes a surprising, checkable subject for a channel — *what if a Category 6 hurricane hit your city?*, *a 2,300-tonne ship engine turns slower than your car at a red light* — and checks it against everything the channel has already said. An agent (Claude Fable 5.2) then writes the episode in one of two pictures. **As code**: a WebGL2 scene with real physics, true scale, a moving camera and a narration script, built with the channel's own renderer kit, style rules and the lessons learned from its past failures. Or **in Blender**: a photoreal 3D scene assembled from free, open-licence models, textures and skies, with real street maps where a place matters. A GPU box renders it; a Chirp 3 HD voice reads the script; a speech recogniser on the same box hears the narration back and times every word, so the captions are the voice, word for word. A critic model, an entertainment reviewer and a series of mechanical gates check facts, readability, pacing, units, the caption sync and the story of the 90-second edit. Only then does it enter the channel's library, ready for its slot.

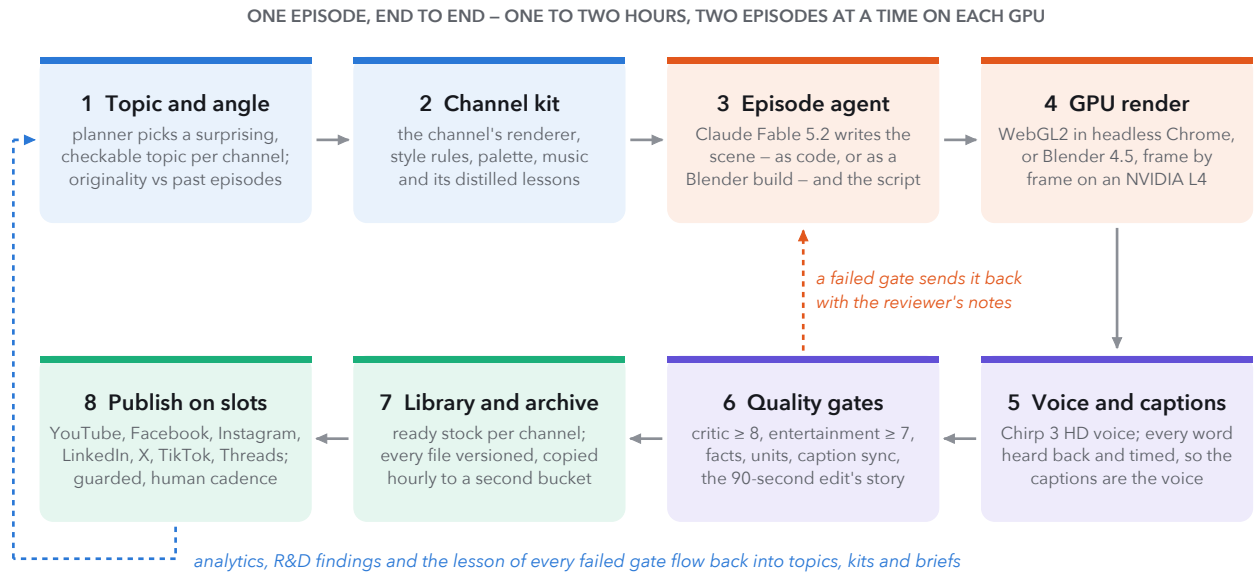


Figure 2. How one episode is made. Every stage is automated; a failed gate sends the episode back with the reviewer's notes, and what is learned flows back into the next briefs.

Two systems, one engine

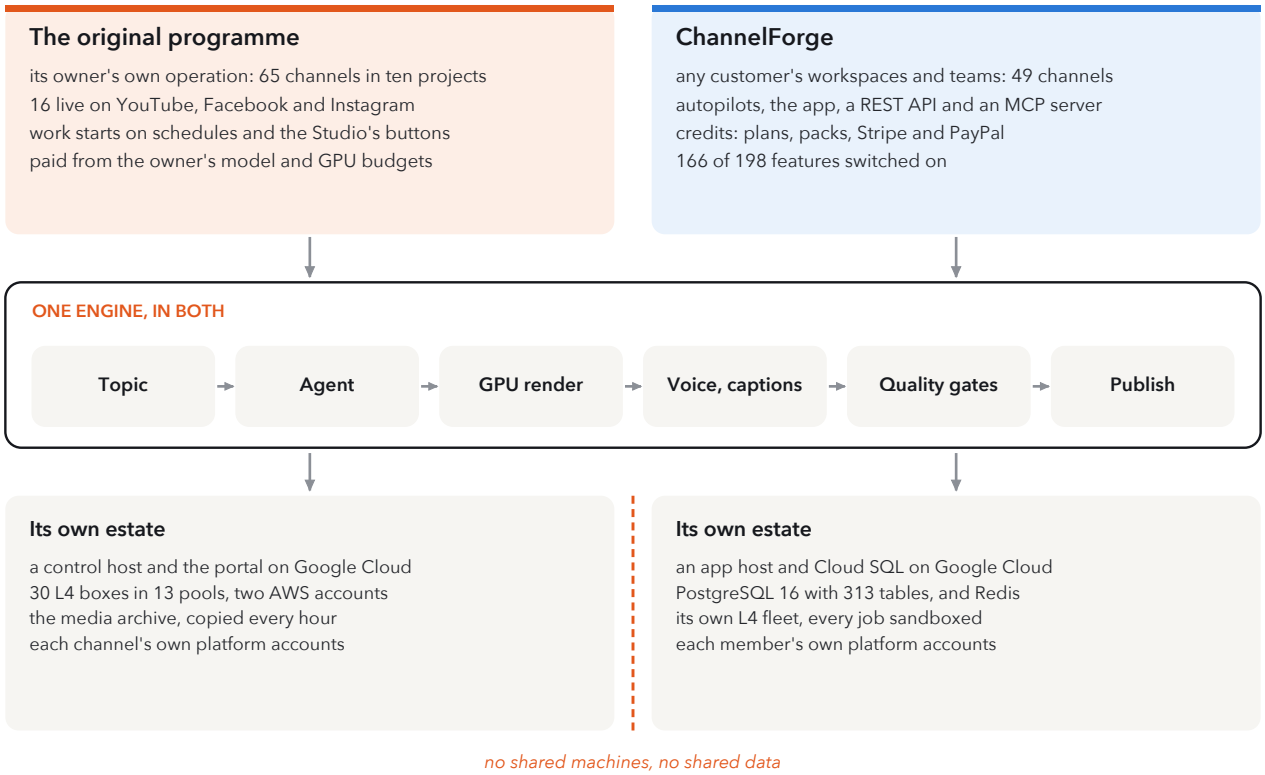


Figure 3. Two systems, one engine: the same pipeline in two estates that share no machine and no data, so either can be sold on its own.

	Original programme	ChannelForge
Who it serves	Its owner: 65 channels across ten projects	Any customer: workspaces with their own channels (49 today)
Web address	content.storeswiki.com (private)	forge.storeswiki.com (public)
How work starts	Schedules (n8n), a scheduler, the Studio's buttons	Autopilots, schedules, the app, the REST API, MCP
Pictures	code-rendered scenes and Blender scenes	code-rendered scenes; Blender built in, switched on per workspace
Publishing	YouTube, Facebook, Instagram, LinkedIn, X, TikTok, Threads	The same plus Bluesky and Pinterest
Money	The owner's model and GPU budgets per project	Credits: plans, packs, Stripe and PayPal
GPUs	30 NVIDIA L4 boxes in 13 pools, two AWS accounts	Its own L4 fleet, isolated from the original
Data	Postgres (portal), files and ledgers (control host)	PostgreSQL 16 (Cloud SQL, 313 tables), Redis
Media	Google Cloud Storage, versioned, copied hourly to a second bucket	Google Cloud Storage, per workspace

The original programme came first and remains the owner's own operation; ChannelForge was built from a copy of it on 18 October and grew into a product with its own control plane, GPU fleet and database. The two share no machines and no data — deliberately: the original can be sold with its channels, and the SaaS as software, independently.

What makes them different

- **Original pictures, not stock or generative video.** Every frame is computed: from code — simulations, maps, cutaways and true-scale comparisons — or rendered in Blender from free CC0 assets. The result is distinctive, cheap to render, and free of the licensing and likeness risks of stock footage, avatars or text-to-video models; every Blender episode says so in its first seconds: *no AI-generated video*.
- **Captions that are the voice.** Since 13 November every new short shows the narration's own words, each one lit as it is spoken, timed by a speech recogniser that hears the finished soundtrack; a gate measures the sync before the video may be posted.
- **Quality before quantity.** A video that scores below 8/10 with the critic, fails a fact, misreads a unit or runs captions too fast is held, never posted. Films have their own floor per act.
- **Self-healing, with a lever for everything else.** A sentinel checks the whole pipeline every five minutes and acts: it makes up a missed slot from surplus, repairs a failed post with backoff, refills stock that runs low, moves work to another GPU, zone or cloud account, and restarts what died. Whatever it cannot fix alone — an account to reconnect, a platform to top up — appears in a failure history beside a button that fixes or retries it.
- **Human cadence.** What the systems write on the platforms by themselves follows no fixed clock: random sessions in each account's own time zone, uneven gaps and varied shapes, so channels never look like bots.
- **Measured, not guessed.** Every post's views, likes, comments and shares are read back; an R&D wing studies what works per channel and feeds it into the next topics and briefs.

CHAPTER 2 · RESULTS

The numbers so far

The original programme published its first posts on 13 October 2026. By 14 November 2026 it had made 2,015 short episodes (1,310 published, 652 ready in stock or waiting for their channel's launch, 30 held by quality gates, 23 rejected by a review) and 28 long-form films. ChannelForge made 944 episodes for its workspaces in the same period. In the eleven days since the first edition both totals more than doubled.

1,513
YouTube Shorts posted

1,479
Facebook Reels posted

1,311
Instagram Reels posted

118
LinkedIn and TikTok posts

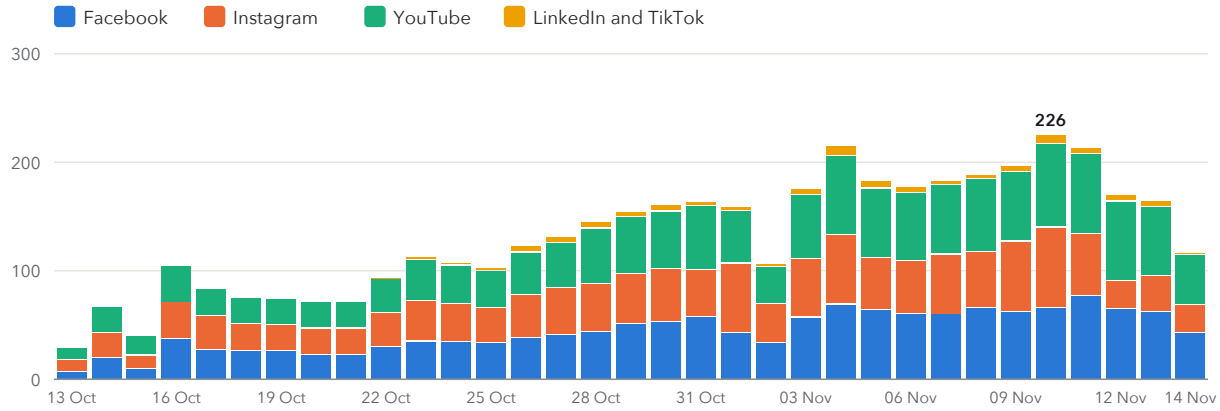


Figure 4. Posts published per day by the original programme, by platform (13 October – 14 November 2026; the last day is partial). The dip on 2 November is the day the whole estate moved to Google Cloud; since 4 November it posts about 190 times a day.

Where the audience is

Across the 16 live channels, the platforms reported 4,434,425 views by 14 November 2026: 1,844,210 on Instagram, 1,555,199 on Facebook and 1,035,016 on YouTube. The accounts have gathered 13,481 followers and subscribers so far, in a month, without any paid promotion — 2.6 times as many as eleven days earlier.



Figure 5. Views by platform, all 16 live channels of the original programme, as reported on 14 November 2026.

Production and cost at a glance

	Episodes	Model cost	Per episode	Avg critic score
Project 1 (4 channels)	536	$5,397	$10.07	8.04
Project 2 (4 channels)	615	$6,141	$9.99	8.15
Studio projects 3–10 (57 channels)	864	$10,266	$11.88	8.06
Original programme, total	2,015	$21,805	$10.82	8.08
ChannelForge (all workspaces)	944	$9,891	$7.02	8.05

Model cost is what the AI models were paid, measured per call, Blender episodes included. ChannelForge's figure covers every job type (channel drafts, kits, brands, insights and so on), its per-episode figure only the episodes. GPU time is in 'Costs and unit economics'.

Output per day

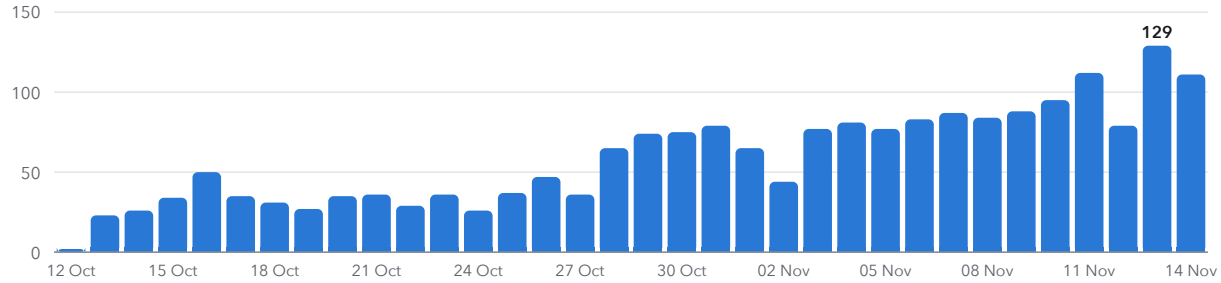


Figure 6. Short episodes made per day by the original programme (12 October – 14 November 2026; the last day is partial). Output roughly doubled from 28 October, mostly in the Studio's projects; from 10 November, when the Blender episodes began, it reached 129 on its busiest day.

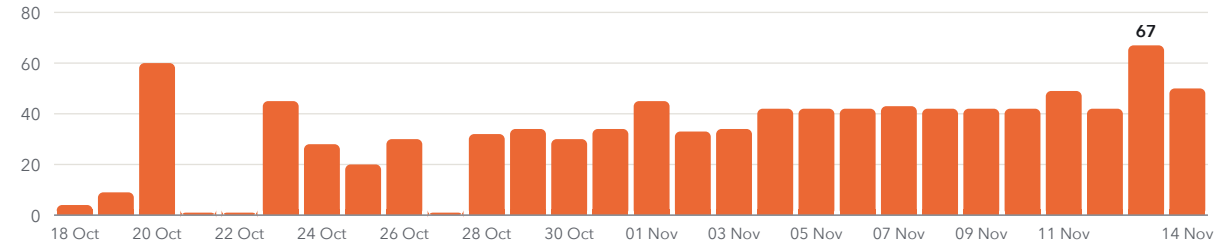


Figure 7. Episodes made per day on ChannelForge, all workspaces, from its launch on 18 October 2026. The peak on 20 October is the browser end-to-end test of the whole product; since 4 November the owner's workspace makes about 42 a day, and more from 13 November as its three new projects began.

PART II

The original programme

65 channels in ten projects, the pipeline that makes their episodes and films — in code and in Blender, captioned word for word — the publishing calendar, every failure and its lever, the self-healing layer, and the portal the owner runs it from.

CHAPTER 3 · THE PROGRAMME

Ten projects, 65 channels

Channels are grouped into projects. A project has its own model budget, its own generation schedule, its own credentials and its own GPU pools, so a problem in one never stops another. Projects 1 and 2 are run from the command line and by schedules; Projects 3 to 10 were created from the portal's Studio, by pressing buttons, and live in a second copy of the pipeline that cannot reach the first two. Projects 7 to 10 — thirty-two channels — were added after the first edition: Project 7 on 5 November, Projects 8, 9 and 10 on 12 November, each with eight niches that overlap no other project's.

Project	Channels	Status	Episodes	Views
Project 1 – science and history	Space Simulated, What if, History Simulated, Body Disclosures	live, 12 slots a day	536	1,192,563
Project 2 – physics and nature	Oddly Simulated, Inside The Machine, Into The Deep, Mind Tricked	live, 12 slots a day	615	2,326,376
Project 3 – mathematics	Quiet Proof	pre-launch, generating every day	16	–
Project 4 – machines and technology	Colossal Scale, Model Cutaway, Downwash, Fracture Atlas, Unseen View, Upright Robots, Signal Voyage, In Phone Charges	live, 24 slots a day	460	915,486
Project 5 – Earth's forces	The Quiet Cavern, Rivers of Ice, Stepped Leader, Millimetre Earth, Slipface, Tephra Town, Tremor Lines, Into the Updraft	pre-launch, generating every day	126	–
Project 6 – the living world	Wild Blueprint, Built to Endure, Deep Time Street, Homeward Atlas, Pursuit Lines, Wild Sensorium, Smaller Than Rain, Swarm Logic	pre-launch, generating every day	116	–
Project 7 – how the built world works	Beforehand, Why That Size, Cofferdam, Below The Block, Domino Trail, Odd Ledger, The Load Path, Throughput Lab	pre-launch, generating every day	60	–
Project 8 – sport, speed and craft	Brake Run, Engineless, Halfspace, Lean Angle, Line Call, Midflight, Railhead, Tenths	pre-launch, generating every day	29	–
Project 9 – making and materials	Endgrain, Heatwork, Within a Micron, Off The Line, Under the Patina, Specimen Bench, Unbuilt, Unfurled	pre-launch, generating every day	28	–
Project 10 – life's strategies	Envenom, Fledge, Host Hijack, Linchpin, Mutual Terms, Mycelium Minute, Outlived, Now You Don't	pre-launch, generating every day	29	–

Pre-launch channels generate every day into a stock of up to 60 waiting episodes (14 in Project 3) – one a day, two in Projects 7 to 10; Projects 8 to 10 add a Blender episode a day and a film every third day as GPU time allows, Project 10's Blender ones once its animals look real – so they launch with weeks of content. Their brand suites, renderer kits and account plans are ready; they go live when their owner connects their accounts.

Every channel is a brand and a renderer

When a channel is created — by hand for Projects 1 and 2, by the Studio's buttons for the others — the system builds three things for it. A **brand suite**: name, positioning, palette, typography, logo lockups, avatars, banners and highlight covers, rendered on a GPU so they match the videos' look. A **renderer kit**: the channel's own JavaScript scene library (its camera grammar, materials, lighting, map and scale tools), verified on the GPU against reference stills. And a **topic bank** seeded with proposals the planner can draw on. The kit is what gives each channel a recognisable visual identity across hundreds of episodes.

> **WHY CODE AND BLENDER INSTEAD OF VIDEO MODELS**
>
> Generative video models produce footage that looks generic, drifts between shots, cannot show true scale or correct physics, and raises likeness and licensing questions. Code-rendered scenes are exact (a 300 km/h blade tip is drawn at 300 km/h), consistent within a channel, cheap to re-render, and entirely original. Blender scenes add photoreal materials, light and real places, built from free CC0 models and street maps – still computed, still exact. The trade is that every scene must be written – which is exactly what the episode agent is for.

The live channels

Each live channel publishes from its own YouTube channel, Facebook Page and Instagram professional account. The handles are the real ones; brand suites (logo, banner, avatar, highlight covers, palette, voice and tone) were generated by the system and applied by the owner. Three channels went live after the first edition — Unseen View, Signal Voyage, Upright Robots — as their owner connected their accounts in the portal; the system took them from there, posting their stock in order, oldest first, and holding back by itself the few early episodes made before the video format's overhaul of 22 October, which have no narration.

Channel	Niche	YouTube	Instagram	Slots	Episodes	Views
Body Disclosures	the human body explained, inside out	@bodydisclosures	@bodydisclosures	2	91	110,209
Space Simulated	space at true scale: journeys, orbits, stars and black holes	@spacesimulated	@spacesimulated1	4	168	449,310
History Simulated	history reconstructed: wonders, engineering and events	@historysimulated1	@historysimulated	3	129	258,041
What if	what-if scenarios: disasters and physics played out	@whatilosophy	@whatilosophy	3	148	375,003
Inside The Machine	how machines work, in 3D exploded views	@machineinsidious	@machineinsidious	3	149	409,135
Mind Tricked	optical illusions and brain teasers (photosensitive-safe)	@mindsparta	@mind.sparta	2	141	212,158
Into The Deep	the deep sea and its life	@deependables	@deependables	3	158	1,323,995
Oddly Simulated	oddly satisfying physics simulations	@oddlysimulations	@oddlysimulations	4	167	381,088
Colossal Scale	machines and structures at colossal scale	@colossalscales	@colossal.scale	3	73	175,660
Downwash	flight and aerodynamics	@downwashphysics	@downwashphysics	3	91	186,255
Fracture Atlas	why materials break: engineering failure	@fractureatlas	@fractureatlas	3	82	394,757
In Phone Charges	the energy budget of everyday devices	@inphonecharges	@inphonecharges	3	76	94,623
Model Cutaway	how AI and modern tech work, in cutaway	@ModelCutaway	@model.cutaway	3	89	53,964
Unseen View	what machines sense that people cannot	@un.seenveew	@unseenview	3	17	1,620
Signal Voyage	how signals travel	@signalvoyager	@signalvoyage	3	16	3,523
Upright Robots	how robots move and balance	@UprightRobots	@uprightrobots	3	16	5,084

Slots: posting times a day, on each of the three platforms; two of each channel's slots are Blender episodes (on the two-slot channels, both). Views are the sum over YouTube, Facebook and Instagram as reported on 14 November 2026. Appendix A lists all 65 channels with every handle.

CHAPTER 4 · THE PIPELINE

How an episode is made

An episode is a folder: a manifest, the scene's code, the narration, the rendered video, its 90-second edit, a thumbnail and a status record. Everything that happens to it is written there, so any episode can be audited, re-rendered or re-posted.

1. Choosing what to make

The topic planner reads the channel's niche, its topic bank, what it has already posted, what performed (from the R&D wing) and a variety rule, and proposes a subject with a hook and an angle: a surprise, a consequence the viewer can see, a number that can be checked. Every topic in the bank carries its angle, the surprising facts, the consequences to show on screen and a 'wow' shot; premises everybody already knows are retired, and so is any topic that would repeat a video already made.

2. Writing the episode

The episode agent — Claude Fable 5.2 since 10 November, working in a sandbox on the GPU box — receives a brief: the topic, the channel's kit documentation, its rulebook (a dozen rules distilled from that channel's past failures), the pacing rules and the quality bar. It writes the scene, the narration script and its on-screen labels, renders stills to look at its own work, fixes what it sees, and iterates until it is satisfied. If a model or a cloud account refuses a call, the call moves down a chain — Fable 5.2, Fable 5.1, Opus 5.5, Opus 5, Opus 4.8 — and across the cloud accounts that may serve it, so a model outage never stops a batch. The scene is either code for the channel's WebGL2 renderer or, for a Blender episode, a Python build of a Blender scene (next chapter).

3. Rendering on a GPU

A code scene is rendered by headless Chrome with WebGL2 on an NVIDIA L4, frame by frame at 1080×1920, and encoded with ffmpeg. Render parity across machines is proven pixel for pixel: the same reference stills come out byte-identical on every box of the fleet. A Blender scene is rendered by Blender 4.5 LTS on the same L4s, headless, in one and a half to five GPU-hours.

4. Narration, captions and the mix

The narration is written for the ear — numbers spoken as words, units as a narrator would say them — and read by a Google Chirp 3 HD voice chosen per channel. A speech recogniser on the GPU then hears every line and times each word, and the captions are drawn from those times: the spoken words themselves, each one lit as it is said (the chapter 'Captions that are the voice'). Music and effects are mixed under the voice, and the result is levelled to streaming loudness (−14 LUFS, true peak below −1 dBTP). A second, shorter cut of 55 to 89 seconds is made for Facebook and Instagram, whose Reels limit is 90 seconds — its joins snapped into the pauses between words, and the edit must tell the story on its own.

5. The quality gates

Gate	What it checks	What happens on a fail
Critic	an independent review of a contact sheet, the captions and the measured motion: visual quality, readability, factual accuracy and policy; 8/10 or more and no blocking issue	revised with the notes, else held
Entertainment	part of the critic's review: would a viewer keep watching — consequences shown, motion, surprise (7/10 or more)	revised, else held
Video and voice	1080×1920 at 30 fps, 120–175 s, H.264 and AAC; the channel's narrator really in the soundtrack	re-rendered, else held
Captions	every spoken word on screen, in order, at the moment it is said; what was drawn matches the word log, frame by frame	drawn again, else held
Caption sync	the finished soundtrack heard again: how far each caption is from its word (median, 95th percentile, share of words matched)	measured on every episode; enforced once its thresholds are final
Text and units	no text outside the one text band; SI symbols on screen, units spoken in words	fixed, else sent back
Picture and page	no black or frozen segment; no JavaScript or shader error in re-rendered stills	re-rendered, else held
Flash	photosensitive safety (WCAG 2.3.1): no more than three flashes in any second	held
Loudness	integrated −15.5 to −12.5 LUFS, true peak at most −0.5 dBFS	re-mixed
The 90-second edit	a viewer model reads the edit's narration alone and must follow it (8/10); else a story editor picks better lines and the edit is cut again	Facebook and Instagram wait for a short version

Episodes the gates hold are kept, never deleted: the owner can review, fix or release them from the portal.

6. Library, archive, publishing

A passing episode enters its channel's library as *ready*. Within minutes the portal's content sync copies every file into the media archive on Google Cloud Storage, where it is versioned and kept, and every hour the whole archive is copied into a second bucket; the control host removes its own copy only once the archive holds the video at its exact size. From the library, slots take the oldest episode each platform has not had, so every platform of a channel posts the same episodes in the same order — a Blender slot takes the oldest Blender episode, a code slot the oldest code one.

Lessons that make it better every day

Every episode that needed a fix, failed a gate or was rejected by a reviewer leaves a lesson. When a channel has gathered a few new ones, a 'doctor' model distils them into that channel's rulebook — at most a dozen sharp rules — which is put at the top of every later brief. Mistakes are therefore made once. The same mechanism carries across channels when a lesson is general.

CHAPTER 5 · NEW SINCE THE FIRST EDITION

Blender episodes: a second kind of picture

Since 10 November the live channels also post **Blender episodes**: photoreal 3D scenes rendered in Blender, the open-source studio tool, from free CC0 models, textures and skies, with real street maps where a place matters. Today they fill two of every live channel's daily slots. 117 have been made so far, in seven of the ten projects, with an average critic score of 8.0.

Why a second picture

Code-rendered scenes are exact and distinctive, but they look like diagrams that move. The bar the owner set for this format was a professional engineering explainer: real materials, real light, real machinery at human scale, labels in the scene and every spoken word captioned. Blender is how a studio would make that, and the system now makes it the same way a studio would — by building a 3D scene, not by asking a model to imagine footage.

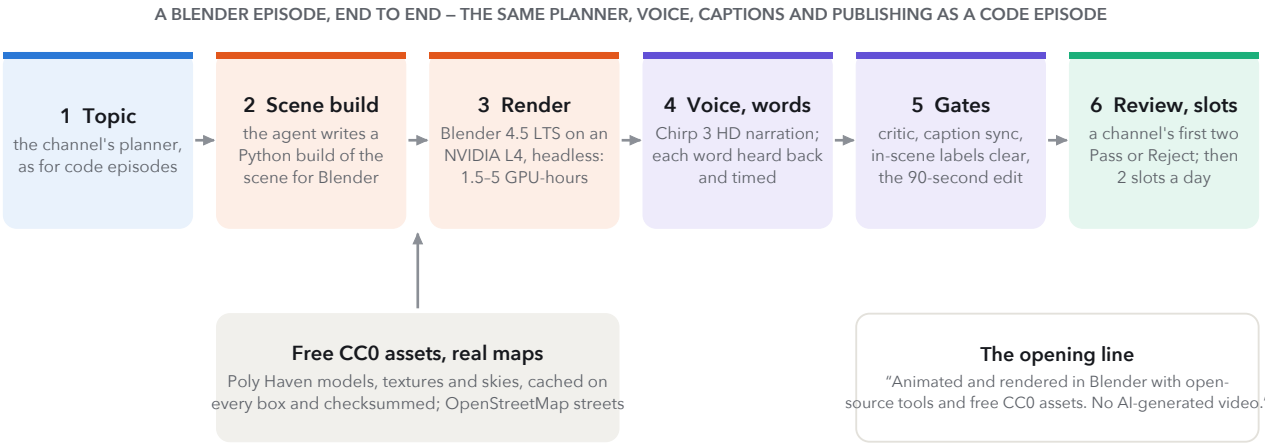


Figure 8. How a Blender episode is made. The topic planner, the narration, the word-timed captions, the review and the publishing are the code episodes' own; what differs is the scene and its render.

How it works

- **The scene is written, then built.** The episode agent writes a Python build of the scene against the system's own Blender toolkit: physically based materials, high-dynamic-range skies, models placed and instanced, a camera language and a colour grade. Nothing is hand-modelled per video. Blender agents run on Claude Opus 5.5: measured on the launch, it scored the same as Fable at a third to a half of the cost.
- **Real assets, legally.** Models, textures and skies come from Poly Haven's CC0 library — free for any use, no attribution owed; a searchable catalogue of 997 skies, 869 textures and 521 models — cached on the boxes and checked by checksum. Where a real place matters, its streets, rivers and rail lines are built from OpenStreetMap.

- **Rendered on the same GPUs.** Blender 4.5 LTS renders headless on an NVIDIA L4 at about 0.8 seconds a frame — an hour of frames per round, one and a half to five GPU-hours an episode with its fix rounds — in Blender pools of their own, one per project, which give way to the live projects' code work when GPUs are scarce.
- **Captions that are the voice.** Every spoken word is on screen, lit as it is said; labels sit in the scene itself, and a check makes sure no label is hidden behind geometry.
- **Nothing toy-like.** The critic holds Blender episodes to a professional explainer's look: since 13 November it blocks anything that reads as a toy — a storm drawn as a disc, a town of white boxes, a few primitives standing in for a machine — and the fix round replaces it.
- **Honest about what it is.** Each Blender episode opens with one line: *Animated and rendered in Blender with open-source tools and free CC0 assets. No AI-generated video.* The script is written with an AI model and the voice is synthesized; the line says exactly what is true of the picture.
- **Reviewed before it is scheduled.** A channel's first two Blender episodes wait for the owner's Pass or Reject in the portal; after that, Blender takes two of the channel's daily slots and code the rest.
- **Every asset is visible.** An episode's page lists every model, texture, sky and map it was built from, with its source and licence.

	Blender episodes	Code episodes
Made	117	1,898
Average critic score	8.01	8.09
Model time per episode	$11.39	$10.79
GPU time per episode	1.5–5 L4-hours, fix rounds included	under an L4-hour, two episodes to a box
Picture	Blender 4.5 LTS: CC0 assets, real maps	WebGL2: the channel's own renderer kit

Model time is measured per call: code episodes since 12 October, Blender episodes since 10 November.

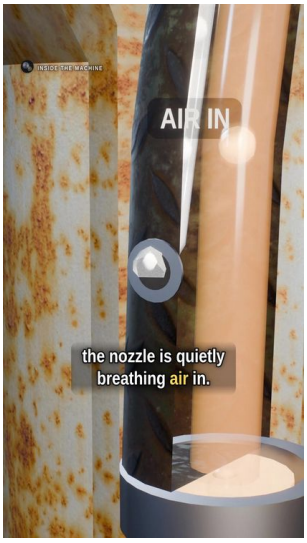
What they look like



Colossal Scale · 29,009 views
This 7 km train has engines in the middle, run by radio



Fracture Atlas · 11,360 views
Why engineers cut half this beam's flange away. On purpose



Inside The Machine · 5,036 views
How does a fuel nozzle know your tank is full? No electronics



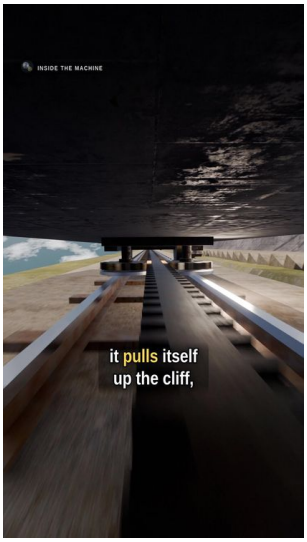
Inside The Machine · 4,763 views
How does a tower crane get taller? It jacks up its own head



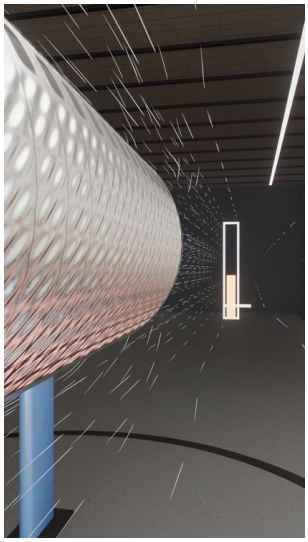
Into The Deep · 3,421 views
3,000 m under the Gulf of Mexico, a volcano erupts asphalt



Into The Deep · 3,111 views
Health-check a whale without touching it? Catch its snot



Inside The Machine · 2,942 views
How does a train climb a 48 % slope? It grips the rail sideways



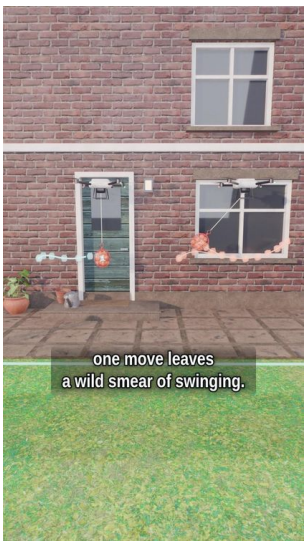
Fracture Atlas · 2,402 views
Why is aircraft skin allowed to wrinkle under load?



Fracture Atlas · 2,355 views
Why are gas cylinders tested with water, not gas?



Into The Deep · 2,253 views
The 8 m 'sea serpent' that swims standing up



Upright Robots · 1,784 views
To stop a swinging load, the drone moves twice

Figure 9. Eleven of the most watched Blender episodes, at most three of a channel, by their real thumbnails: the scene, its in-scene labels and the word being spoken lit in gold (views on their best platform).

CHAPTER 6 · NEW SINCE THE FIRST EDITION

Captions that are the voice

A caption should be what the narrator says, when the narrator says it. Since 13 November that is what every new short of the original programme shows: the narration's own words, two short lines at a time, each word lit as it is spoken — timed not by a guess from the script but by a speech recogniser that hears the finished voice.

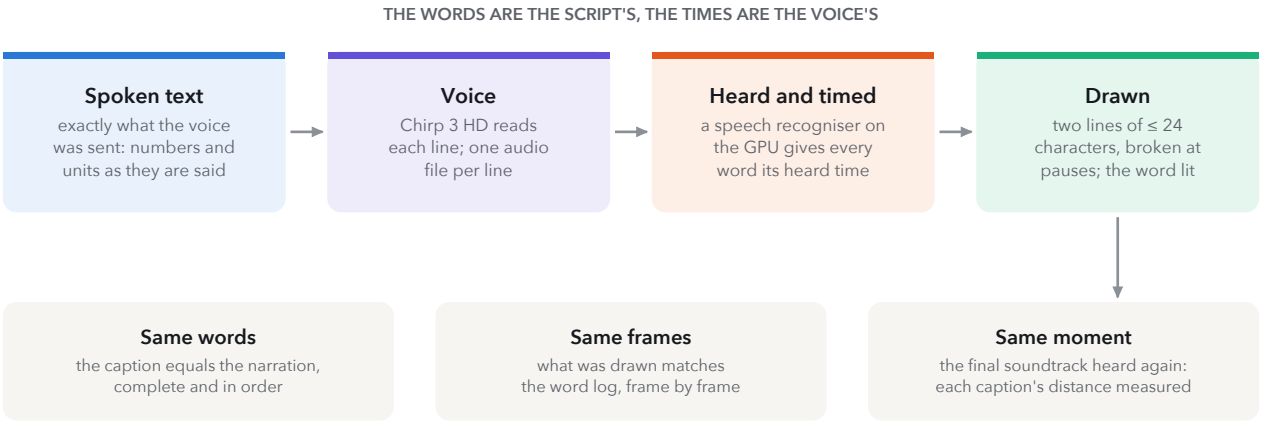


Figure 10. How a caption gets its timing. The words are the script's; the times are the voice's, heard on the GPU box after the narration is made, and checked again on the final soundtrack.

How it works

1. **The words.** Each line's caption is exactly what was sent to the voice — the spoken form, with numbers and units as the narrator says them — so the screen and the speaker never disagree.
2. **The times.** When a line has been voiced, a speech recogniser (faster-whisper large-v3, on the GPU) hears it and returns a time for every word; the script's words take those heard times.
3. **The drawing.** From the first frame, the captions show two lines of at most 24 characters, breaking at the speaker's pauses, the current word lit. The 90-second edits and short versions are cut between words, so a caption is never sliced.
4. **The checks.** The caption text must equal the narration, complete and in order; what was drawn must match the word log frame by frame; and the finished soundtrack is heard once more to measure how far each caption is from its word.

How close it is

Episode	Words	Matched	Median	95th percentile
Project 1, What if – the first captioned episode of Projects 1 and 2 (13 Nov)	265	98.1 %	0.01 s	0.21 s
Project 2, Oddly Simulated – its first (14 Nov)	279	96.1 %	0.01 s	0.11 s
Project 10, Fledge – a Studio sample (13 Nov)	–	98.8 %	0.01 s	0.23 s
A Blender episode – the bar this format set	–	99.3 %	0.01 s	0.09 s

Distances between a caption and its spoken word, measured on the final soundtrack. Across Project 1's first captioned batch – eight episodes in all four of its kits – 95 to 99.6 % of words matched, with the 95th percentile between 0.09 and 0.21 seconds.

Where it stands

The Studio's projects went first, after test videos of each kind, on 12–13 November; Projects 1 and 2 followed on the evening of 13 November, and every short they have made since carries word-timed captions, its 90-second edit included. Short versions, translations and the long-form films are next — the films' captions are built and wait for their test film — and the sync gate, measured on every episode today, becomes a hard gate when its thresholds are final. ChannelForge's version is built and merged, not yet deployed, behind feature flags, to be switched on for customers after a trial on the owner's workspace.

CHAPTER 7 · FILMS

Long-form films

Beside the shorts, the programme makes long-form films: narrated documentaries in eight acts, about nine minutes long (two calm 'journeys' run thirty), made by the same engine at a higher bar. 28 have been made so far (25 approved), and 71 film posts have gone out on YouTube, Facebook and Instagram.

The film autopilot

Every day at 00:30 Rome the autopilot picks the next channel in a rotation across the live channels, plans a film for it — research, a fact-checked outline, acts with their own hooks — and makes it on a dedicated GPU box. Each act is reviewed on its own; an act under the bar gets an improvement round, and a film that is still under its floor is held for the owner instead of being posted. Films that pass are narrated, then posted the next day at 20:00 Rome with platform-specific metadata, chapters and a Reels teaser — and a film that is late is posted as soon as it is ready.

Since 12 November the Studio's newest projects make a film for each channel every third day as GPU time allows; the first, Host Hijack's, waits for its channel's accounts.

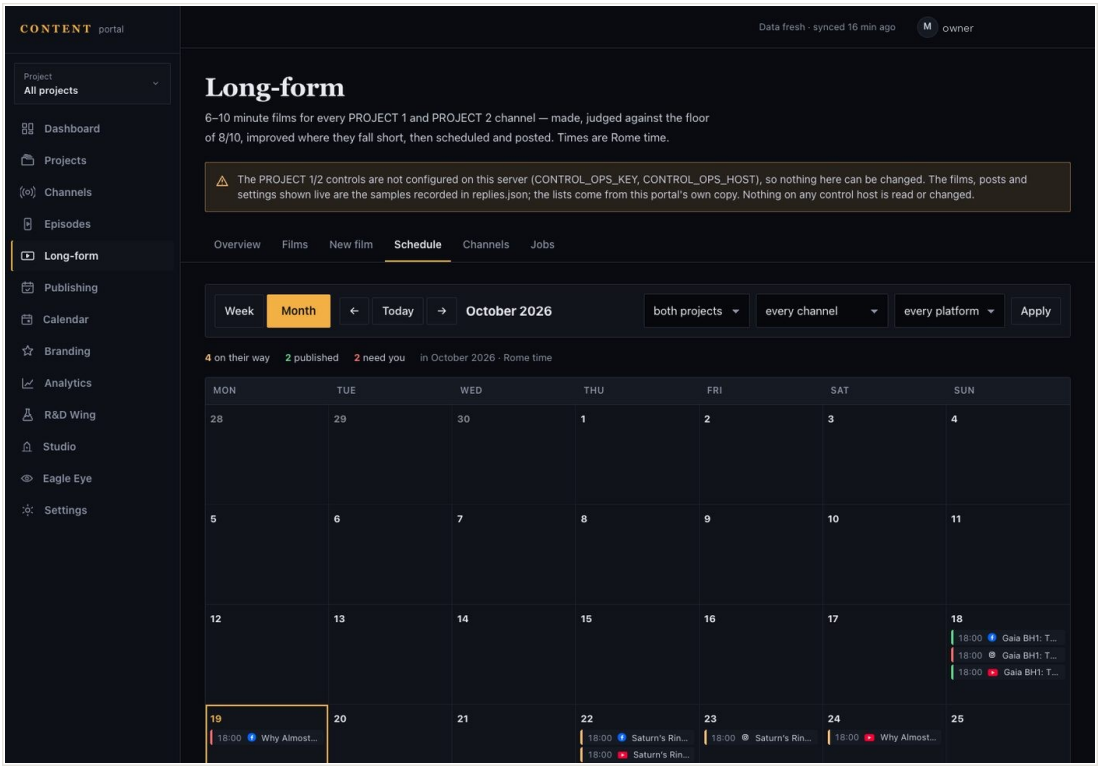


Figure 11. The portal's long-form calendar: each day's film, its channel and its state (development build, demonstration data).

When a film is held

Films are where the bar is hardest to meet: a reviewer watching a nine-minute film finds what no short shows — a readout that drifts from the narration, a label that contradicts a fact, a line that describes a shot the act never drew. Between 9 and 13 November five days' films were held in a row, and the way they were rescued is now part of the system:

- **Fix the picture by hand, never by an AI rework.** The act's own code is patched — a label, a readout, a timing — and only the picture is rendered again; every measured number on screen stays as it was.
- **Fix the words where they came from.** The narration is written from each act's script beats; a beat that describes a shot the act never drew is rewritten to what is on screen, and the film doctor, a model that reads the reviewers' notes, does the same by itself.
- **A floor that is fair.** Since 14 November a film is released at 8/10 overall when each act scores at least 7 with its facts and policy confirmed and nothing blocking — four of the six films held before it had a single act at 7.

What a film costs and earns

A film takes a GPU box for several hours and costs about $156 in model time on average, because every act is written, rendered, narrated and reviewed like a set of shorts; the autopilot caps a film's spending, revisions included, and has a kill switch. In return a film reaches long-form audiences, who watch for longer, and gives its channel the watch time the platforms reward.

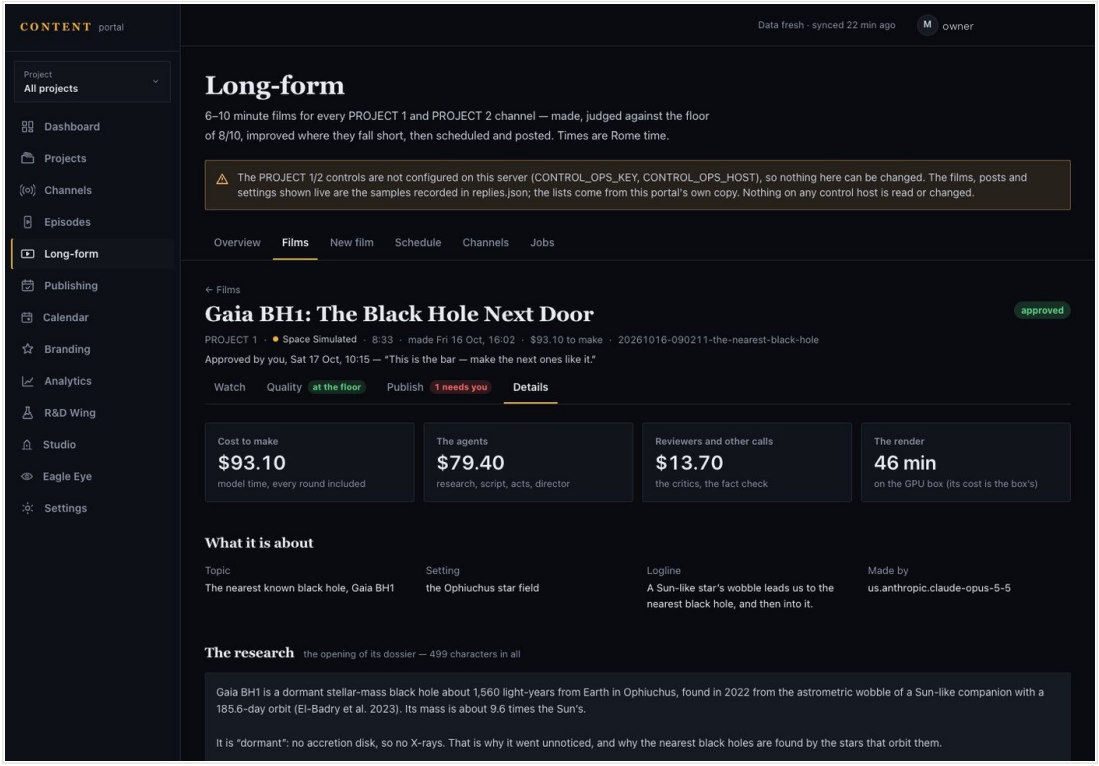


Figure 12. A long-form film's page in the portal: what it cost to make, its agents, reviewers and render, and its research (development build, demonstration data).

CHAPTER 8 · PUBLISHING

Publishing and the calendar

A channel's calendar is a list of slots — two, three or four a day at the hours its audience is awake — on each platform it posts to. The 16 live channels post 48 episodes a day across three platforms, 144 platform posts a day — two of each channel's slots carrying Blender episodes — plus a curated LinkedIn queue and mirrors to X and TikTok.

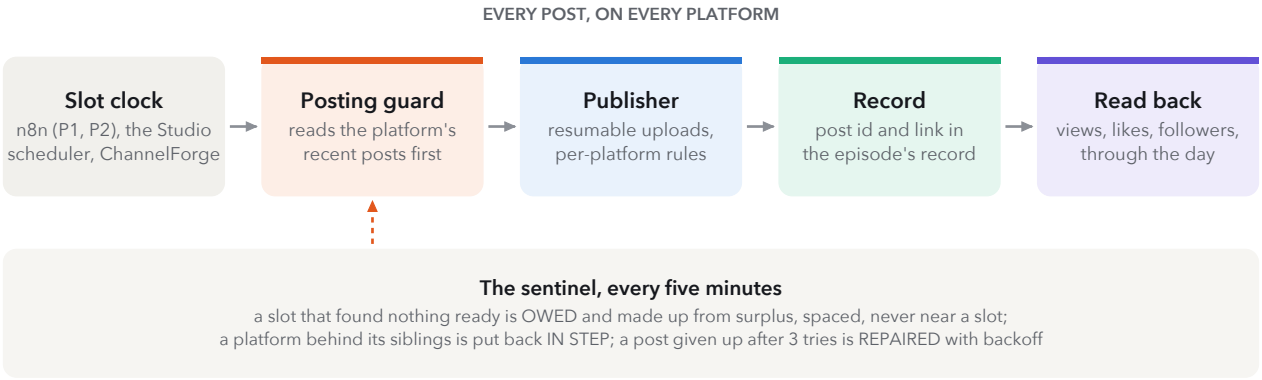


Figure 13. Every post passes the same steps. The posting guard reads what the platform already holds before anything is sent, so a restored or duplicated host can never post an episode twice.

Platforms and how they are used

Platform	What is posted	How
YouTube	the full episode as a Short; long-form films	Data API, resumable uploads; a watch re-uploads any video YouTube leaves stuck in processing
Facebook	the ≤ 89 s edit as a Reel, with an explicit thumbnail	Pages API (Reels), per-Page tokens
Instagram	the ≤ 89 s edit as a Reel	Instagram API with Facebook Login; since 13 November Meta fetches the video from a 12-hour signed link (Projects 1, 2 and 4), uploads as the fallback
LinkedIn	a curated queue: 3 posts a weekday, 1 on weekend days, 4 h apart	the owner's profile, channels take turns
X	mirrors what went out on the three: the full episode when it fits X's 140 s, else the edit	API v2; waits whenever X's pay-per-use credits run out
TikTok	three new episodes a day per channel, through ChannelForge's per-post screen	Content Posting API; posts stay private ('Only me') until TikTok's audit passes
Threads	mirrors, per channel when switched on	Threads API; apps published, switched on per channel

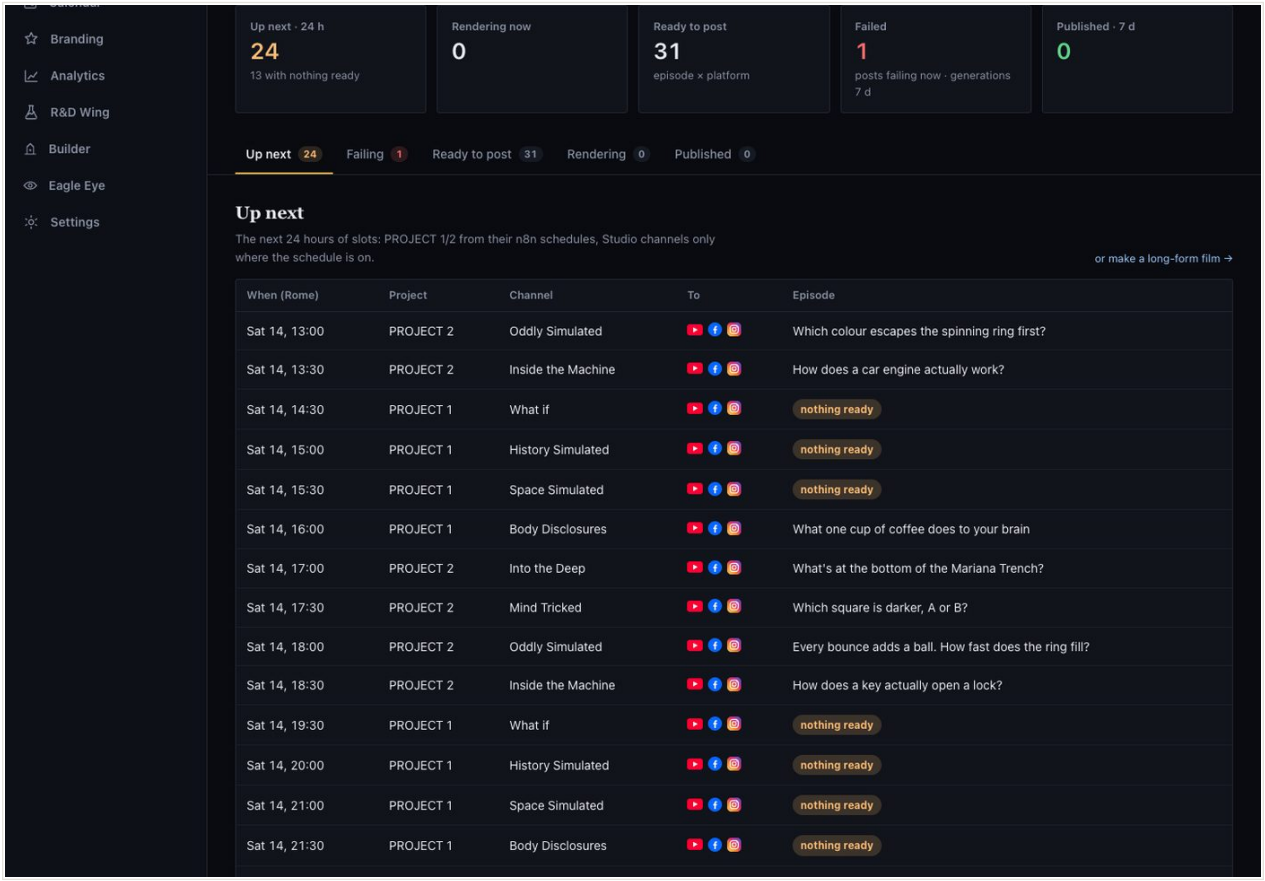


Figure 14. The portal's publishing view: the next 24 hours of slots across the projects, and what is failing, ready, rendering and published, each in its own tab (development build, demonstration data).

Looking human

Platforms flag accounts that behave like machines, so nothing here behaves like one. Titles, captions and hashtags are written per episode; make-up posts are spaced well apart and never land near a slot; and the inbox assistant answers comments on the channels' own posts in the channel's voice — minutes to hours later, in random sessions in the account's own time zone, with varied lengths and no stock phrases, and never every comment.

Schedules

Projects 1 and 2 run on n8n: one workflow per channel holds its slot clock, and one per project starts its generation twice a day (07:00 and 19:00 UTC for Project 1, 13:00 and 01:00 UTC for Project 2). Each calls the control host over SSH with a key whose forced command can only run the control program, never a shell. The Studio's projects use the Studio scheduler, which keeps each live channel's stock ahead of its slots within the project's daily budget, now with videos per day set per channel and slots only on the platforms a channel is really connected to. Daylight-saving folds were found to skip timers, so every interval timer now runs in UTC.

A channel goes live by itself

Launching a channel is now one task for its owner: create its accounts, connect them on the portal's Connections page and tick its slots. The system checks each connection with a test read, takes the channel's stock in order — oldest first — and leaves out by itself any episode made before the video format's overhaul of 22 October, which has no voice. Signal Voyage, Unseen View and Upright Robots went live this way on 13 and 14 November.

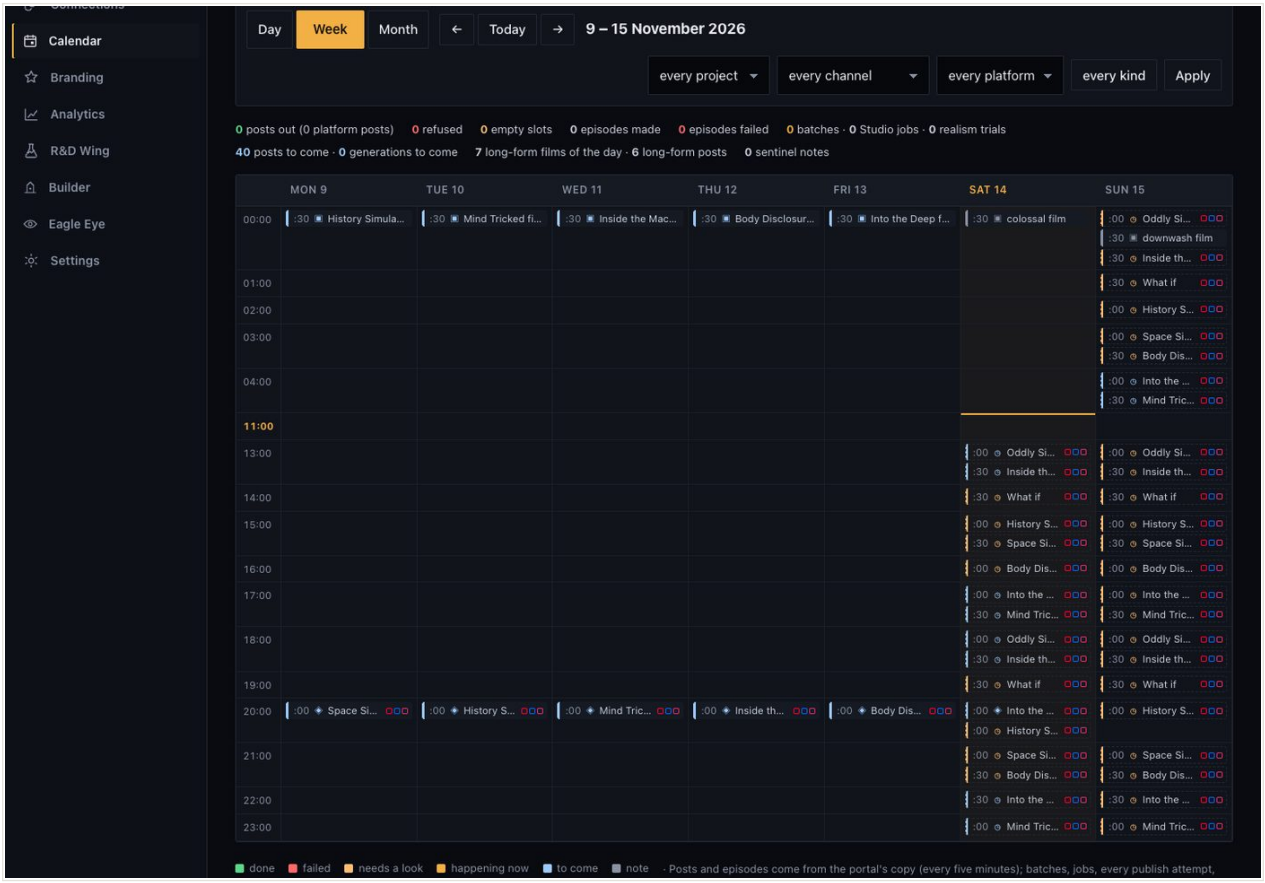


Figure 15. The calendar's week: every slot of every project, each with the episode it will post, the films of the rotation and what went out (development build, demonstration data).

CHAPTER 9 · NEW SINCE THE FIRST EDITION

Every failure has a lever

A platform refusing a post is a normal day's work: a token expires, an account needs a check, a service times out, a pay-per-use balance runs out. What matters is that nothing is lost, that every attempt is on the record, and that whatever needs a person is in front of them with the button that fixes it — in the same place the problem is shown.

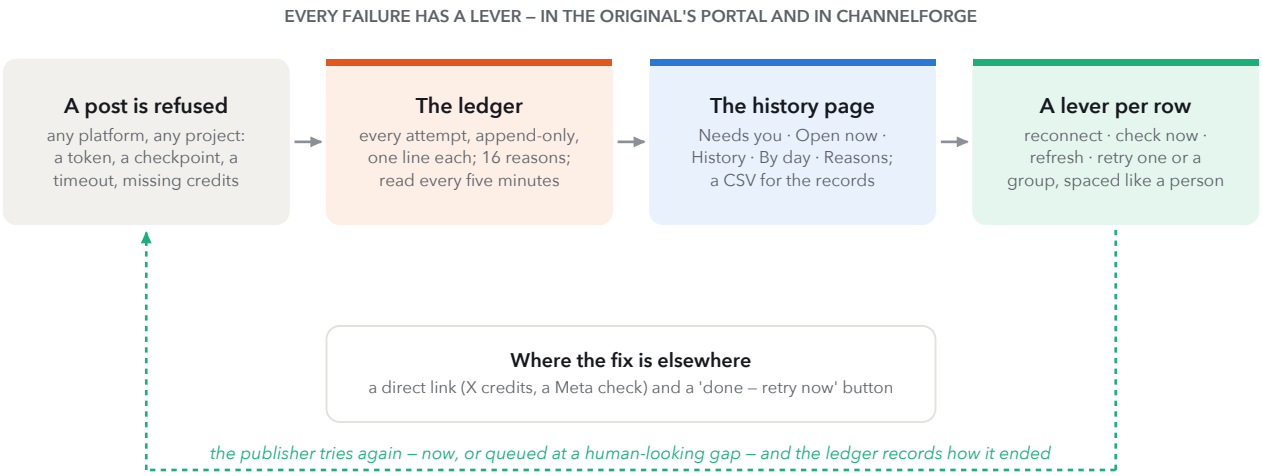


Figure 16. A failed post's way back. The publisher keeps retrying on its own; what it cannot fix appears on one page with a lever beside it.

The ledger

Since 13 November every failed publishing attempt, in every project and on every platform, is written to an append-only ledger: one line per attempt, with the platform's own words, the account, the episode, a reason from a fixed list of sixteen, and how it ended. Backfilled from the records, the history goes back to 25 October: 1,029 failed attempts on 521 posts, of which 490 are resolved. 721 of the attempts are X waiting for its API credits; of the rest, the biggest kind is Instagram saying it could not process a video, which led to the fetch-first upload described in the previous chapter.

Reason	Failed attempts	What it means
X credits	721	X's pay-per-use balance is used up: posts wait, nothing is lost
Media processing	119	the platform could not process the video; most go out on a retry
Account unlinked	81	an Instagram account was unlinked from its Page: the owner relinks it
Timeout or network	36	the platform did not answer in time: retried with backoff
Meta checkpoint	10	Meta asks the account's owner to confirm it in the app
Other	62	tokens, rate limits, posts that may already be live

Every failed publishing attempt of the original programme, 25 October – 14 November 2026, by reason. One attempt is one line: a retry that the platform refuses again is a second attempt, never a duplicate.

The page, and the levers

The portal's Failures page reads the ledger every five minutes and shows it five ways: what **needs you** (grouped by account), what is **open now**, the full **history** with filters and a CSV, **by day** and **by reason**. Every row and every group has a lever: *Reconnect* the account through its sign-in, *Check now*, refresh a Meta or LinkedIn token, or *retry* one post or a whole group — spaced like a person would, stopped at the platforms' daily caps. Where the fix can only happen elsewhere — topping up X, confirming an account in Meta's app — the lever is a direct link and a 'done, retry now' button. The morning report e-mailed at 08:00 has a section on the last day's failures.

ChannelForge has the same, for every workspace: a Failures tab under Publishing for its members, and an admin page for the operator, with the same levers and the same rule — a group retry is spaced, capped and never posts in a burst.

CHAPTER 10 · OPERATIONS

Self-healing operations

The owner's rule is simple: no slot is ever missed and no episode fails for ever. A sentinel enforces it. Every five minutes it observes every part of the pipeline, acts on what is wrong and records why.

Check	What it does when something is wrong
Slots	a slot n8n never fired is published on the spot; a slot that found nothing ready becomes an owed post
Stock (pace)	if a slot would go empty before the next scheduled generation, it starts a top-up batch
Owed posts	made up from surplus stock only: never two posts on a channel within 110 minutes, never within 55 minutes of a slot
In step	a platform behind its siblings is caught up, so every platform posts the same episodes
Repair	a post given up after three failures is retried at 2, 4, 8… hours, up to ten times
Catch-up generation	missed videos are made in gaps before the next scheduled run, charged to a budget
GPUs	a box running with nothing asking for it, a reconciler that stopped, a start refused for hours
Cloud accounts	an account that refuses to start a box is passed over like a full zone, and the pool's next box is started – in the other account
Yield	pre-launch work gives its GPUs to a live project whose batch has waited eight minutes for a box
Batches	a batch silent for ten minutes is marked failed, so the next one can start
YouTube	a video YouTube leaves stuck in processing is uploaded again by itself
Archive	a video kept on the host because the archive lacks it, a stale archive listing
Posting guard	a platform that could not be checked before posting; a host whose records are behind
Host	disk space, load, the portal's pulls, n8n's reachability and errors

Every incident is written to a log the portal shows as Eagle Eye's Observer, and errors are emailed. A failure the sentinel cannot fix by itself — a platform token the owner must renew, credits to top up at a provider — is said in plain words with what to do, and since 13 November with the lever that does it (the previous chapter).

A REAL TEST, ON 14 NOVEMBER

At 03:49 UTC one of the two AWS accounts stopped accepting new GPU starts, and the pools kept retrying the refused box while the other account's boxes stood idle; one film batch failed waiting. By 06:40 the GPU agent passed a refused account over like a full zone: the long-form film being re-rendered moved to the other account in twelve seconds, and every pool now starts its boxes there. The three Blender pools whose only box was in the stopped account are held – set to release themselves – and their slots post code episodes meanwhile, so the channels went on posting. The incident became lesson fourteen.

Lessons, written down

Every incident that cost something — a slot, a batch, a day of films — ends as a numbered lesson in the operations handbook, with what happened, why, and the rule that now prevents it: fourteen so far, from 'capacity before volume' and 'a retry must never erase the record of a failure' to 'one blocked cloud account must never stall a pool that has other boxes'. The same lessons are turned into the code's own checks.

Eagle Eye

Eagle Eye is the read-only automation console: the pipeline as stages (schedules, generation, render, review, publish), every channel's state today, what is going wrong grouped by cause, and everything in time order. It reads the control host through a forced command that can only report — it cannot change anything.

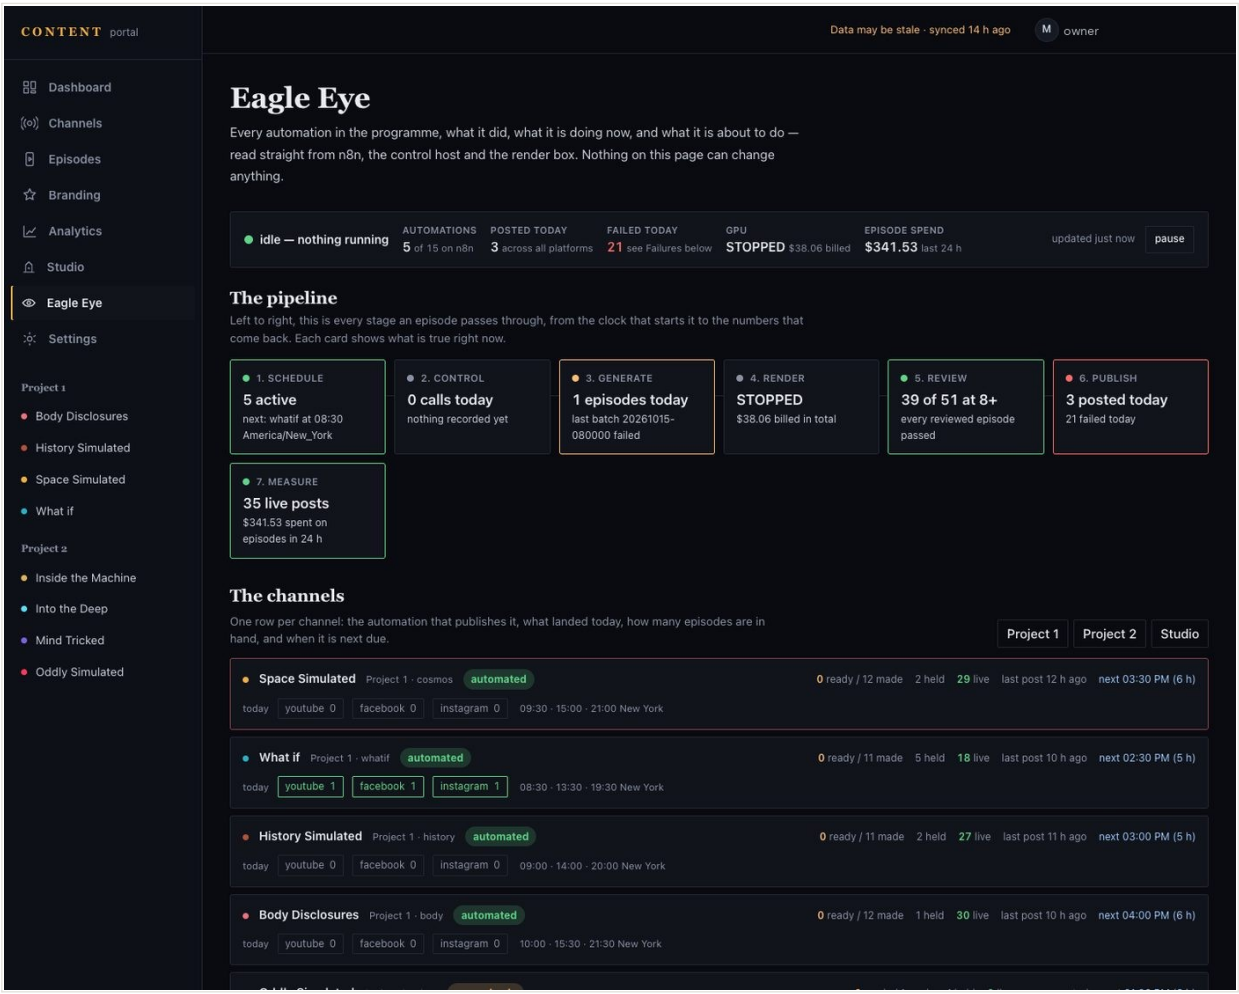


Figure 17. Eagle Eye, the automation console (development build, demonstration data).

CHAPTER 11 · THE OWNER'S COCKPIT

The content portal

content.storeswiki.com is where the owner runs the programme: a Next.js application with its own Postgres database, two-factor sign-in and rate limiting, which reads the control host over SSH every five minutes.

Section	What it is for
Dashboard	today across all projects: the newest episodes, what needs attention, views, likes and live posts
Channels and episodes	every channel in tabs (overview, episodes, brand and words); every episode with its video, scores, captions word by word, Blender assets, facts, cost and posts; retry, hold, release
Studio and Builder	create projects, find niches, draft channels, build kits and brands, generate code or Blender episodes, schedule videos per day per channel – by pressing buttons
Publishing and calendar	every slot of every channel, what went out and what is owed, in tabs; the week and month
Failures	every publishing failure with its lever (the chapter 'Every failure has a lever')
Connections	every account of every project with its health and levers, and which account each channel posts as
Long-form	films, their acts and quality, the rotation, approvals and publishing
Analytics	views, likes, comments, shares and followers per post, channel and platform, over time
R&D wing	measured intelligence per video: what works, findings, forecasts, playbooks
Eagle Eye	the automation console (the chapter 'Self-healing operations')
Settings	sessions, two-factor authentication, recovery codes

Every page in tabs

A portal that grows with a programme grows 'God pages' — pages that show everything at once and take a minute to read. On 13 November the 24 such pages among the portal's 58 were split into tabs, sections and drawers, in four lanes released one at a time; every tab has its own address, so a link opens exactly the view it names. 'Goldens' taken before the move — 105 states of the 58 pages, each recording every heading, field, action and link — prove that nothing was lost. The same rules were applied to ChannelForge first.

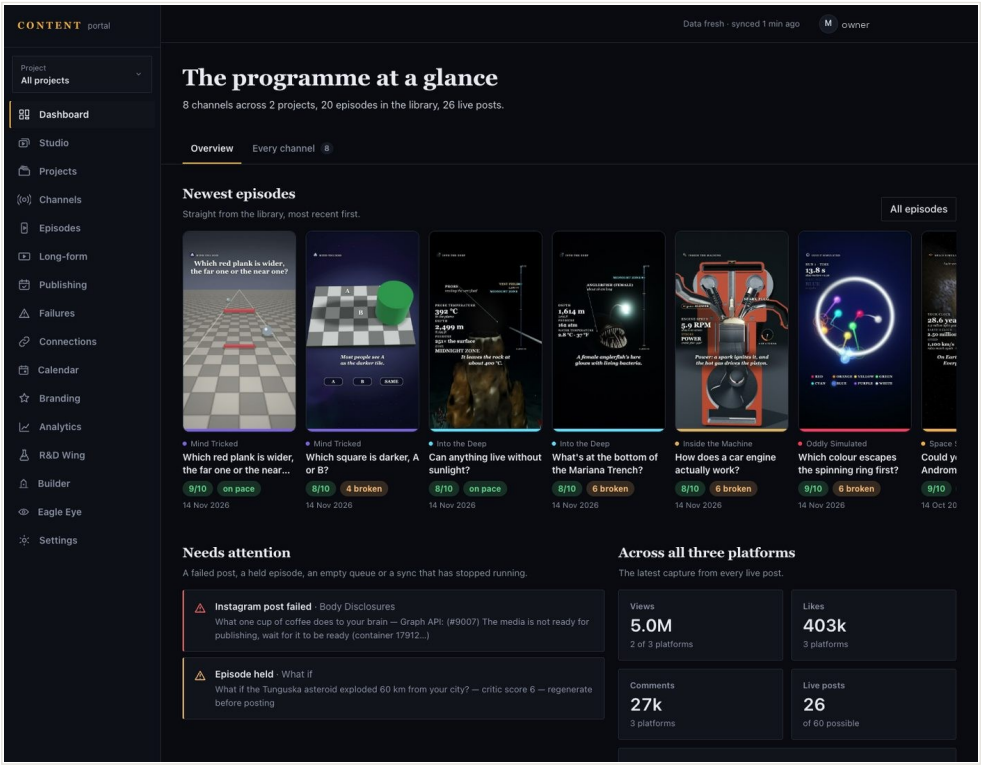


Figure 18. The dashboard: the newest episodes with their scores, what needs attention, and the three platforms' numbers (development build, demonstration data).

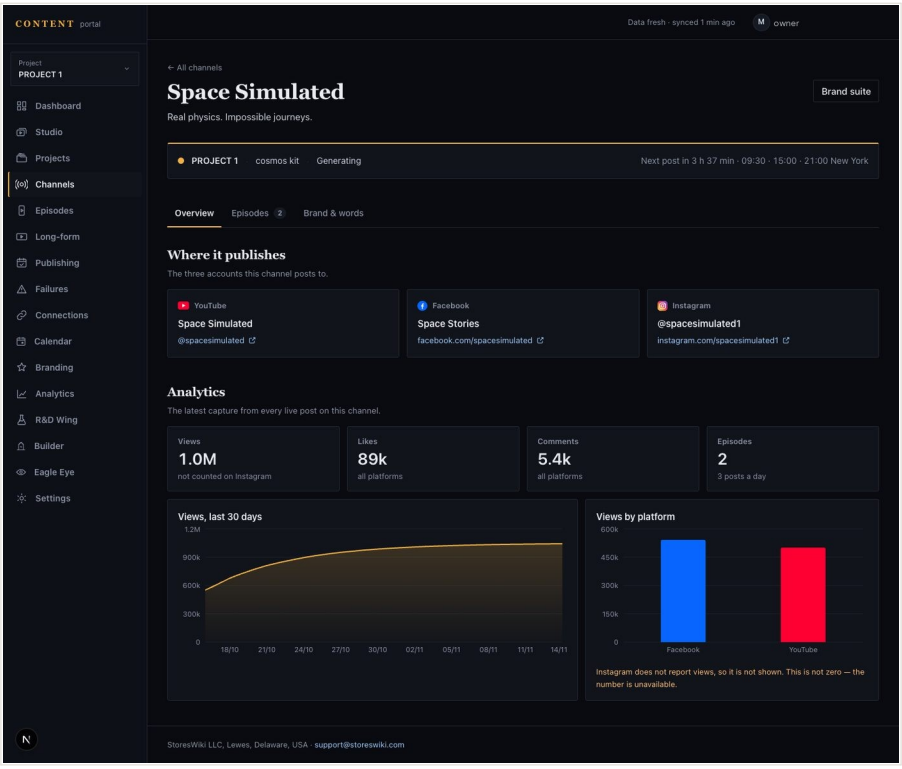


Figure 19. A channel in tabs — overview, episodes, brand and words — with its three accounts and its numbers (development build, demonstration data).

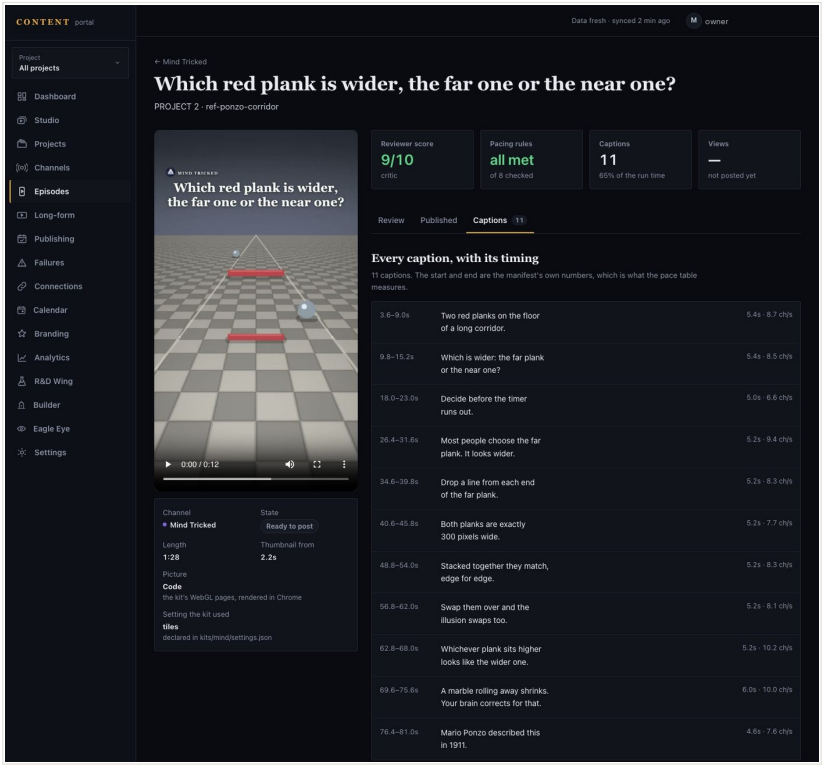


Figure 20. An episode's captions tab: every caption with its timing and pace, beside the video and its scores (development build, demonstration data).

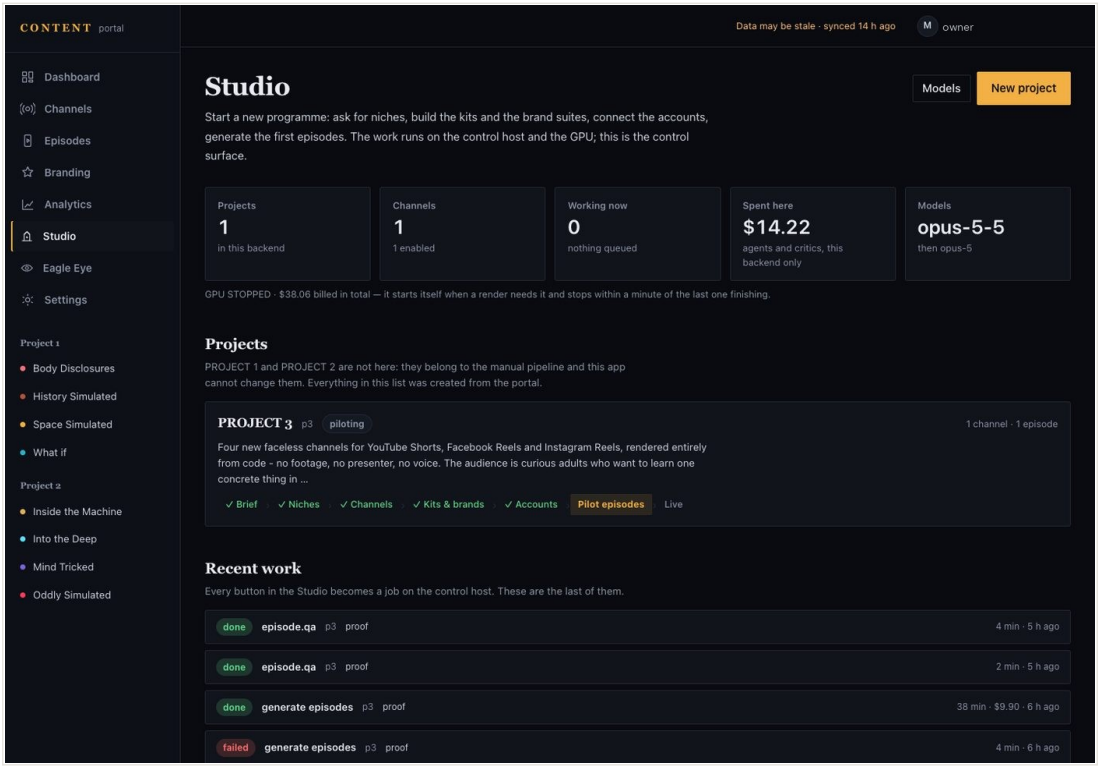


Figure 21. The Studio: projects, niches, channels and jobs, driven by buttons (development build).

PART III

ChannelForge

The same engine as a multi-tenant product: what customers get, how it is built, how tenants are kept apart, how GPUs are shared fairly, how it charges, and how other software drives it.

CHAPTER 12 · CHANNELFORGE

The product

ChannelForge (forge.storeswiki.com) lets anyone run faceless channels the way the original programme does. A customer describes what they want; ChannelForge researches niches, proposes channels, builds each one's brand and renderer, plans topics, makes episodes on its GPUs, checks them with a critic and publishes them on schedule — and then tells the customer what worked.

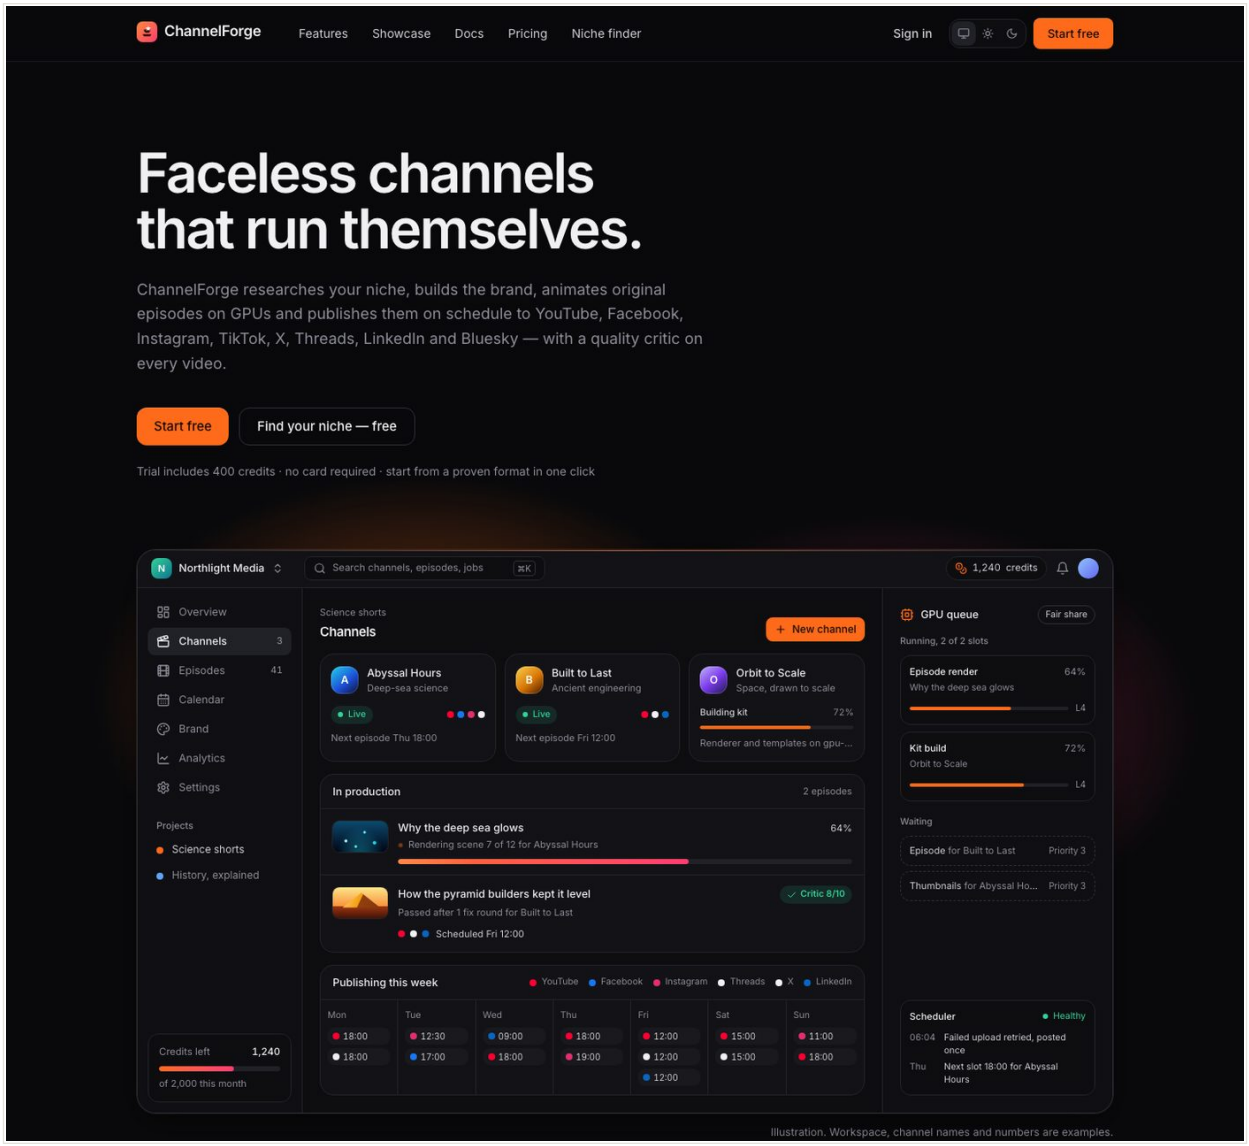


Figure 22. ChannelForge's home page, captured live on 14 November 2026.

The customer's journey

1. **Find a niche.** A free niche finder suggests channels with demand, competition and monetisation notes.

2. **Draft a channel.** Name, positioning, voice, topics and posting times, ready to edit — or start from a proven format in one click.
3. **Build it.** A kit (the channel's renderer) and a brand suite are built and verified on a GPU.
4. **Make episodes.** One at a time, in batches, or on an autopilot that keeps a buffer ready.
5. **Approve or auto-post.** Review links, approval queues, or straight to the schedule.
6. **Publish everywhere.** YouTube, Facebook, Instagram, LinkedIn, X, TikTok, Threads, Bluesky and Pinterest, from the customer's own accounts.
7. **Learn.** Insights, retention maps, outlier detection and a coach turn results into the next briefs.

The features every customer has

A weekly coach brief per channel	First-episode picks	Outlier radar
Agency mode: client workspaces, white label	GPU job isolation	Packaging lab: titles and thumbnails
An episode editor with revisions	Hook lab: alternative openings	Pinterest publishing
Bluesky publishing	Inbox assistant (automatic replies)	Product channels and kits
Brand voice, with consent	Insights and findings	Profit per channel
Breakout detection and follow-ups	Language twins (translations)	Retention learning
Caption styles	Link-in-bio pages	Seamless loop endings
Carousels (Instagram, LinkedIn, Threads)	Long-form films	Storyboard approval
Channel mascots	MCP server for AI agents	Studio: comments, replies, catalogue
Clipping long videos into shorts	Media kits for sponsors	Threads publishing
Comment themes and topics	Monetisation tracker	TikTok publishing
Fifty languages	Originality guard	Voice casts

The 36 features open to every workspace on 14 November 2026, each behind a flag. Another 130 are switched on in the owner's own workspaces, where every new feature is proven before it opens to customers (the next chapter); 32 more are built and wait for a partner, a key or the launch.

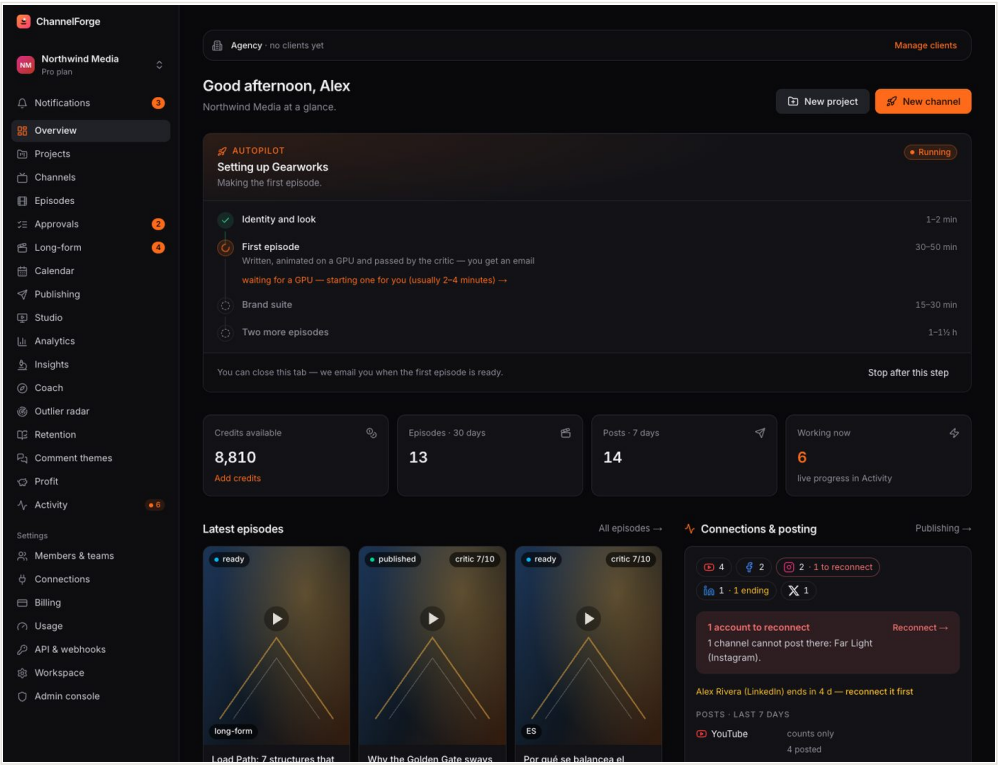


Figure 23. A workspace's overview: the autopilot setting up a channel, credits, episodes and activity (development build, demonstration data).

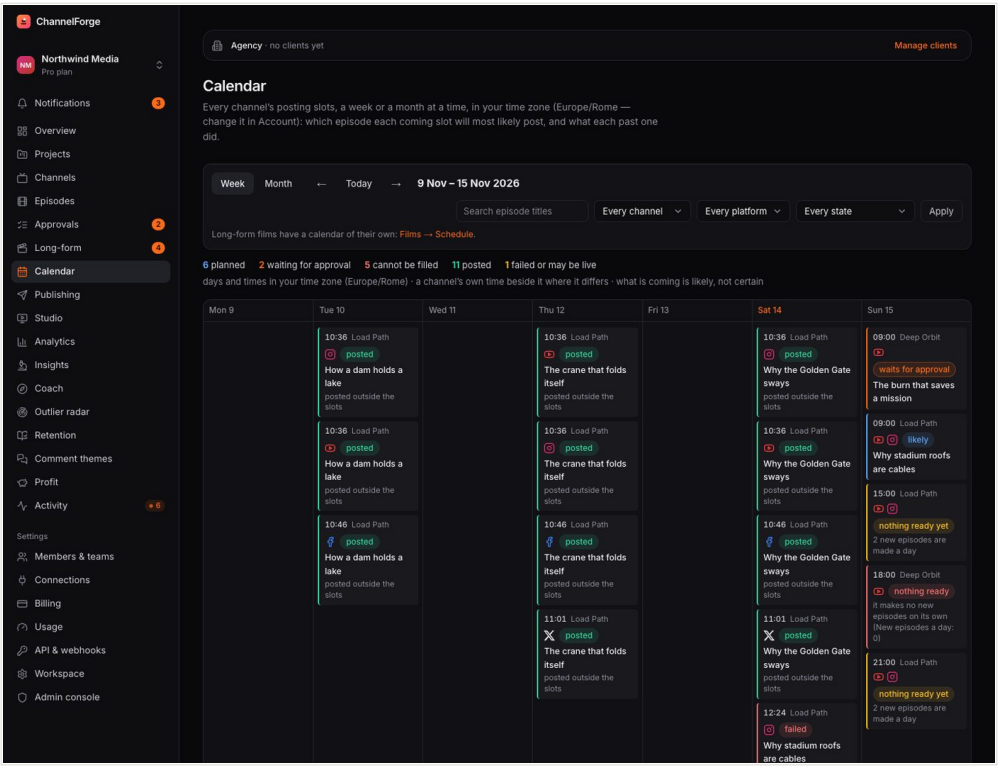


Figure 24. The publishing calendar across channels and platforms (development build, demonstration data).

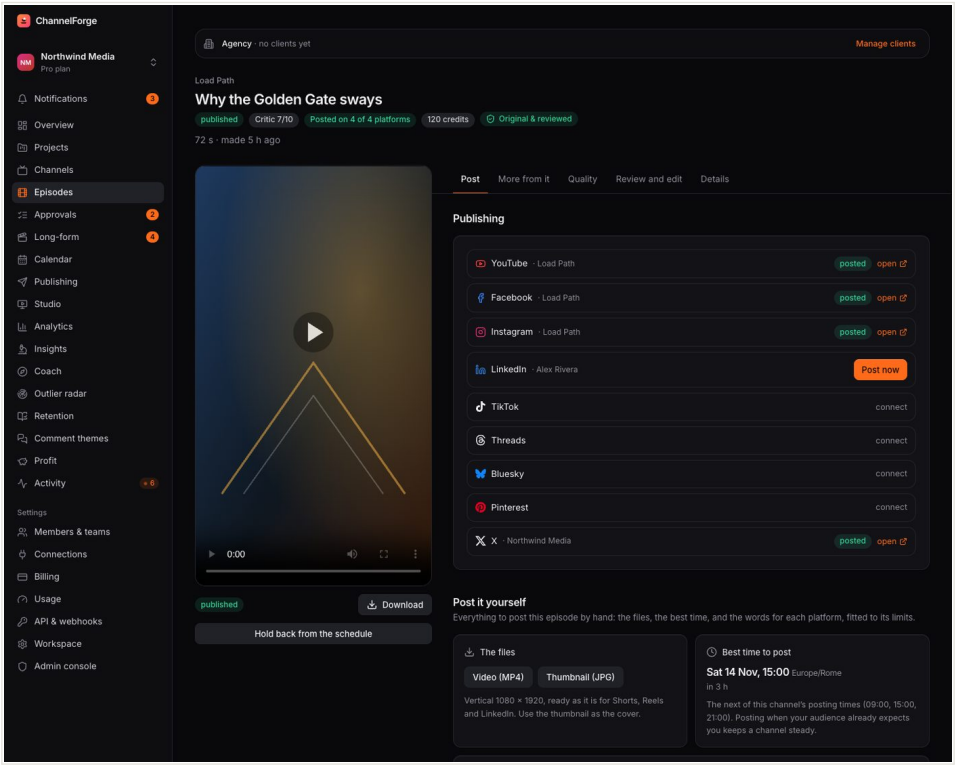


Figure 25. An episode in tabs — post, more from it, quality, review and edit, details — with its post on every platform and everything needed to post it by hand (development build, demonstration data).

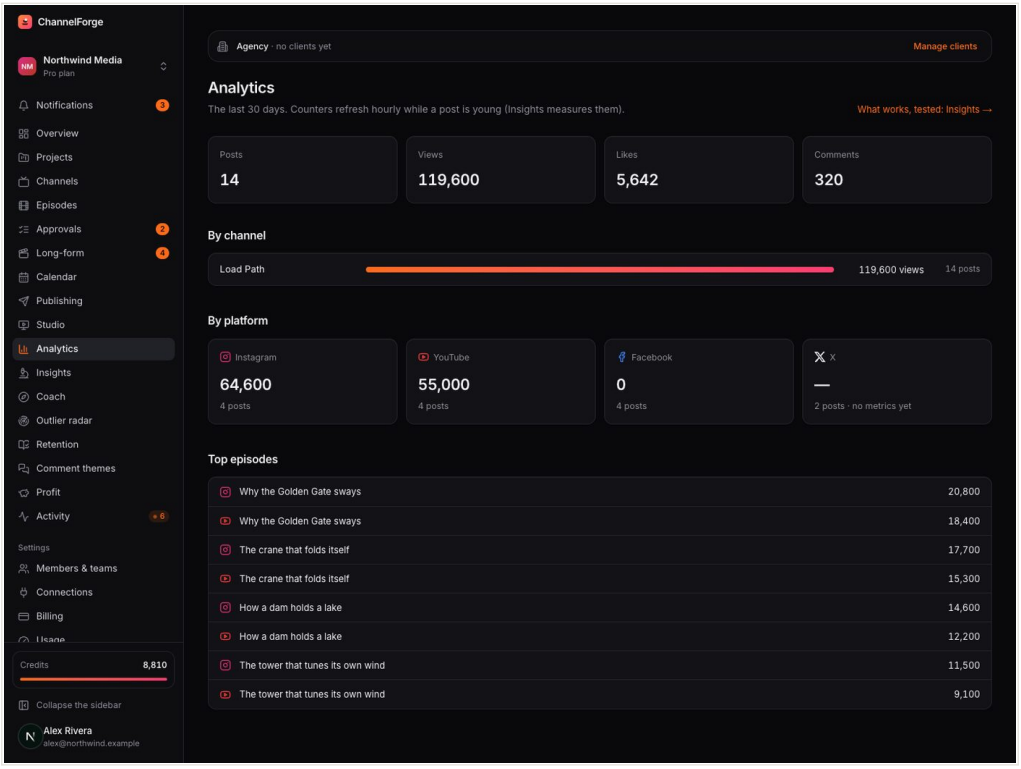


Figure 26. Analytics: posts, views, likes and comments by channel and platform, and the top episodes (development build, demonstration data).

CHAPTER 13 · NEW SINCE THE FIRST EDITION

What ChannelForge gained since the first edition

In the eleven days after the first edition ChannelForge received 873 commits — 2.7 times as many as in its first sixteen days. Most came from a method: study a mature product in the same market feature by feature, decide what ChannelForge should have, build it in parallel lanes behind feature flags, prove it in the owner's own workspaces, and only then open it to customers. 130 such features are switched on in the owner's workspaces today, on top of the 36 every customer has.

Three studies, about 160 features

Product studied	Features read	Built in ChannelForge	Released
A social-media scheduler and composer	1,356	59 backlog items and a wave of ~70 smaller ones; 78 flags	5–6 November
An AI content-creation suite	1,629	85 backlog items; 65 flags	7 November
The owner's own YouTube tool	624	7 fixes and 17 ranked features; 17 flags	8 November

No code was copied: each feature was redesigned for ChannelForge's stack, after checking what it already had. Some things were declined on purpose – a free-form composer, AI image and video generators, synthetic presenters, stock media, review gating, lifetime plans.

What a member can do now

- **Write in the channel's voice.** Channel words and style — avoid lists, preferred terms, calls to action, writing samples, brand facts — fed to every writer; topics imported from a CSV, a pasted list or RSS feeds; a channel's own documents used as sourced facts; pronunciations; an article written from any episode or film.
- **Check before it ships.** Hand-edited text checked for limits, AI wording, numbers and units; a second-opinion critic that watches and hears the finished video (only a lower score holds it); films' dated and numeric facts re-checked against Google Search; library search by meaning, and an originality check that catches paraphrases.
- **Package and publish more.** Thumbnails touched up with a live preview inside each platform's safe area, covers chosen from a frame, Mastodon, first comments through the official APIs, tracked short links, Pinterest Pins from carousel slides, brand plates for every platform, a post taken down from the app, connection health with a notice for every expired account, posting days, windows and pinned slots.
- **Run the inbox as a team.** Assignments, private notes, labels, approved answers, digests by e-mail, Slack or Discord; a review studio for YouTube that checks a pack before it goes, compares the live description, moderates up to fifty comments in one call and keeps series playlists; YouTube's own search terms and retention against similar videos.

- **Work as an agency.** Client roles, approval with a note, master downloads, a monthly report in PDF, a support desk whose tickets are tied to the job, episode or film they are about, and answers drawn only from ChannelForge's own documentation.
- **Ask the product.** 'Ask ChannelForge', an in-app assistant, works through the member's own MCP tools and proposes what it would do with its price; nothing runs until the member clicks.

Every page in tabs

ChannelForge's 139 pages were audited on 7 November; the 24 that had grown into 'God pages' — the channel page alone ran to about seventeen screens — were split into tabs on 8 November: the channel page into eight (overview, topics, look and voice, schedule, distribution, growth, earn, settings), the episode page into five. Every tab has its own address, old links still land in the right place, and a test of 147 page states proves nothing was lost. Its public pages now read the live feature flags, so the site never promises a feature that is not open.

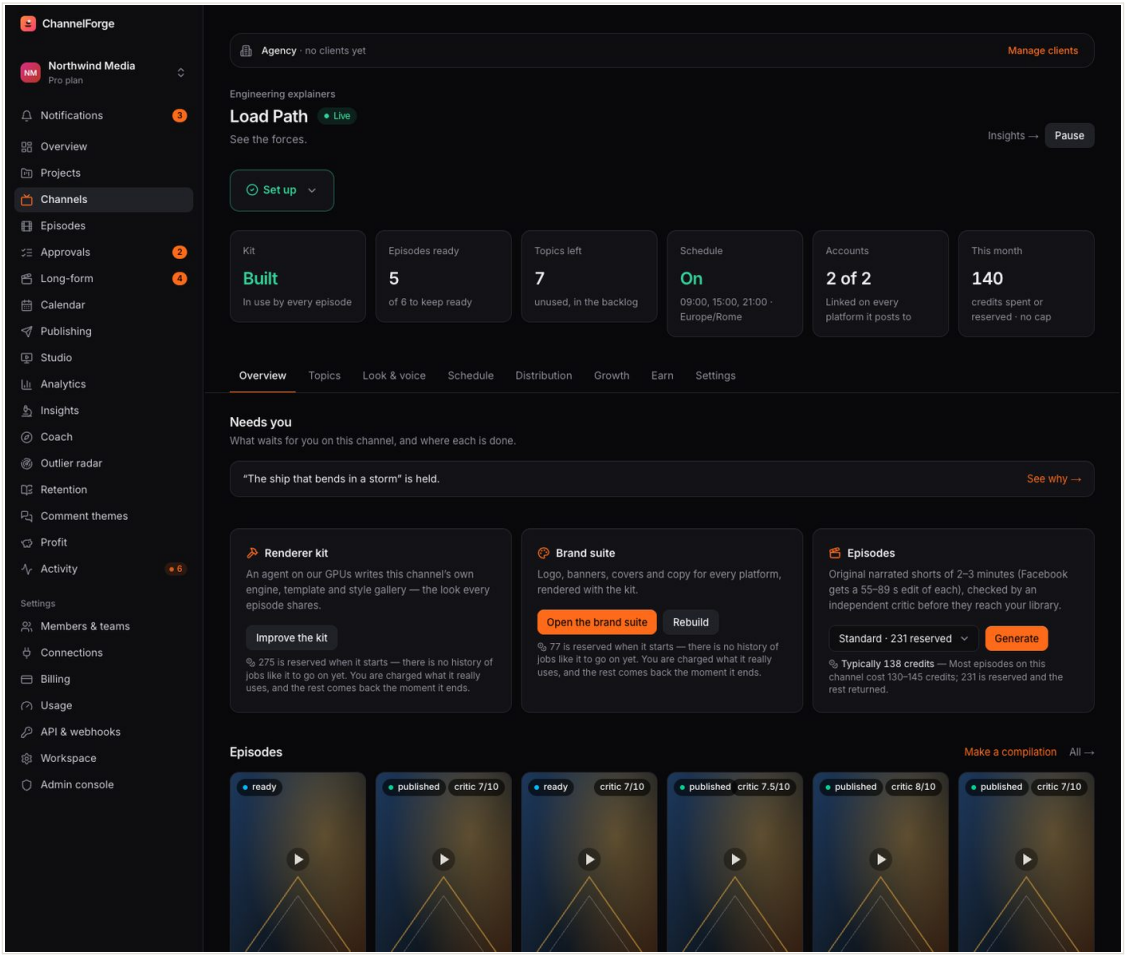


Figure 27. A channel in eight tabs: its kit, brand and episodes, what needs the member, and what to press (development build, demonstration data).

Blender and word-timed captions, for members

Blender episodes reached ChannelForge on 10 November, on the same terms as the original's: a choice of picture per channel and per episode, every word captioned, the opening line on screen, every asset credited on the episode's Quality tab. A GPU box sets itself up for Blender in about twelve minutes with 5.5 GB of cached skies, textures and models, because a job's box has no internet access. A Blender episode is a price class of its own, charged at cost times the plan's markup — about 700 to 1,100 credits on Pro — and never charged when it is held. It is switched on in the owner's workspaces first, where the critic holds it to the same 'nothing toy-like' bar. The word-timed captions are built, reviewed and merged behind two flags, not yet deployed; they go live after a trial.

Every failure has a lever, here too

ChannelForge's publishing failure history went live on 13 November: one row per failed attempt, sixteen reasons in plain words, the same five views as the original's, a morning digest — and, beside every row, the button that fixes it: connect again, retry now, retry the waiting posts (spaced like a person, at most twenty a press and three on YouTube), skip, or 'it is there'.

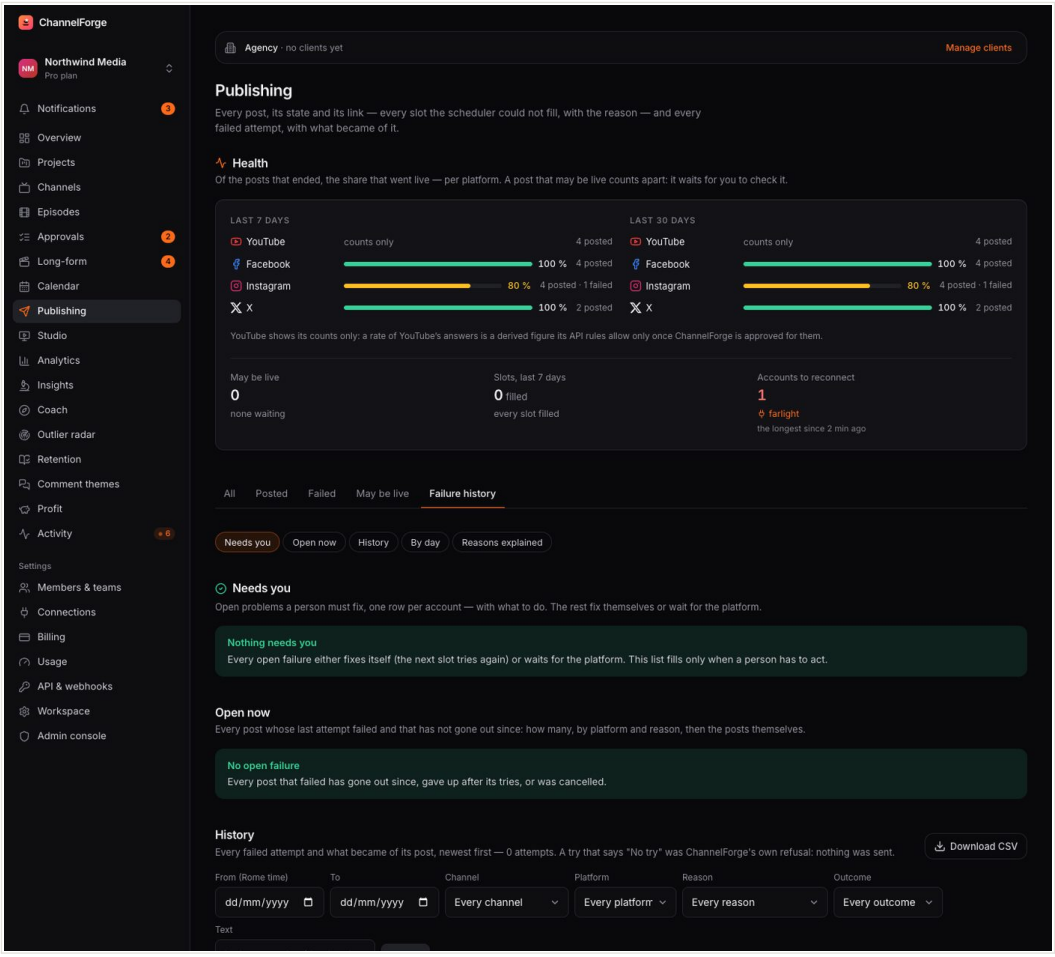


Figure 28. Publishing › Failure history: each platform's health, then what needs the member, what is open, the full history and the reasons explained (development build, demonstration data).

The numbers against the first edition

	3 November	14 November
Feature flags	36, all open	198: 36 open, 130 in the owner's workspaces, 32 off
Database tables	142	313
Publishing platforms	9	11 — Mastodon and Telegram added (Telegram waits for its bot)
Job kinds	54	103
App pages (staff pages)	130 (14)	200 (42)
REST API route files · MCP tools	32 · 20	41 · 30
Public documentation pages	75	82
Plans and credit packs	5 and 4	the same, at the same prices

Counted from ChannelForge's code and, for the flags, its production database. Webhook events are unchanged at 23. Built but waiting for the launch of selling: receipts, promotions, local-currency prices, add-ons and automatic top-ups.

CHAPTER 14 · CHANNELFORGE

Architecture

ChannelForge is a web application, a control plane and a GPU fleet. The web application is what people and programs talk to; the control plane does the work; the fleet renders. They meet in a PostgreSQL database and a Redis instance.

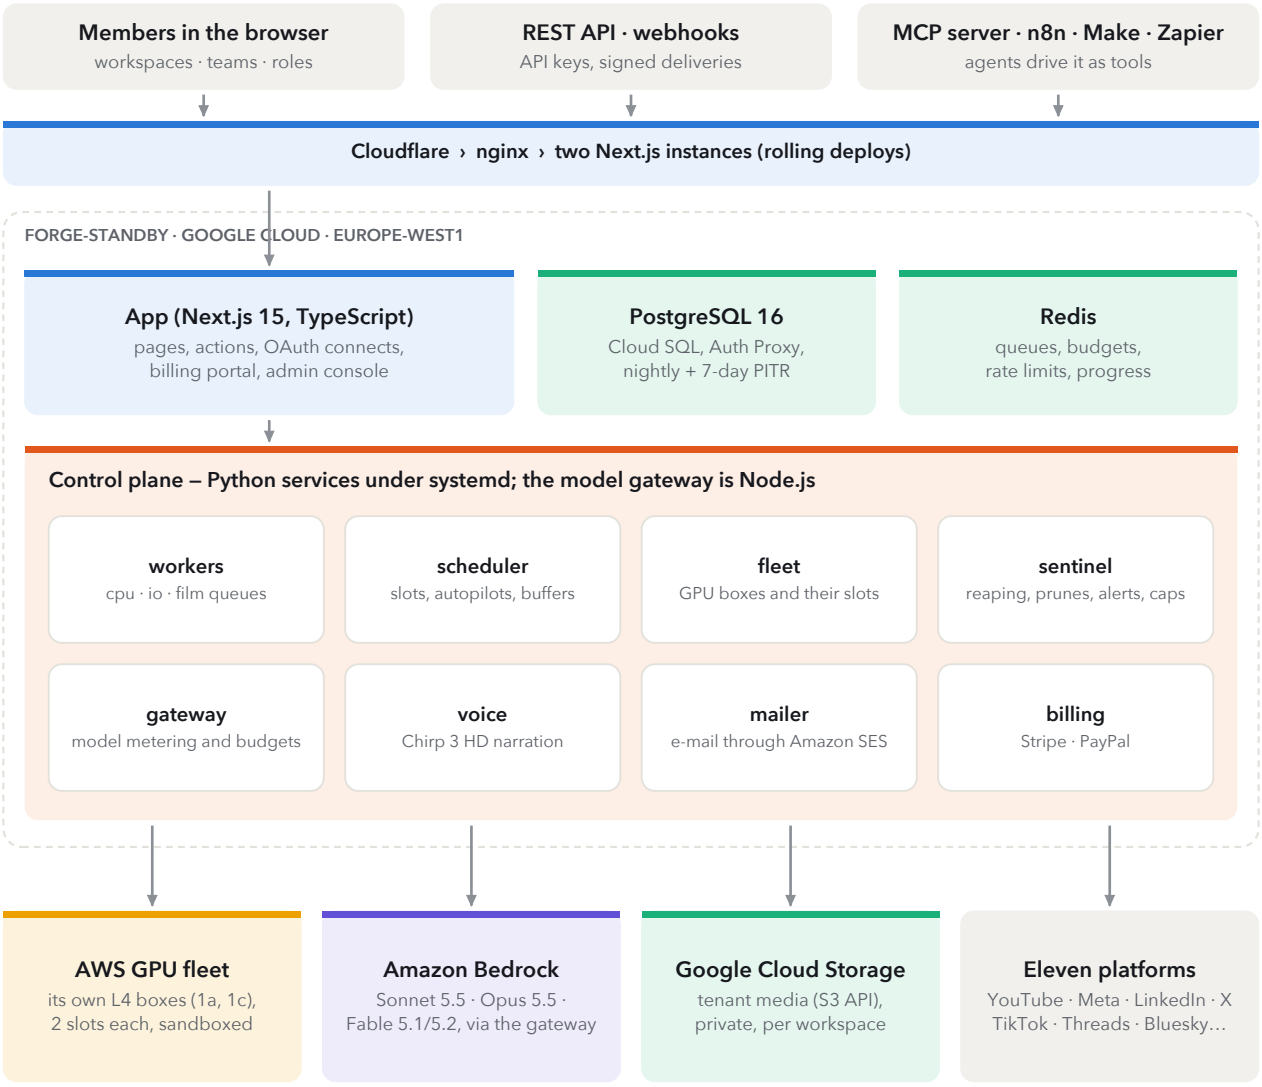


Figure 29. ChannelForge's architecture.

The web application

A Next.js 15 application in TypeScript (React 19 server components, server actions, Tailwind CSS) serves the marketing site, the documentation (82 pages), the signed-in product and the admin console (42 staff pages). Two instances run behind nginx, which retries a request on the other instance, so a deploy rolls one instance at a time; a release builds first and stops the control units around its migrations. Prisma maps the database's 313 tables.

The control plane

Service	Role
worker-cpu / worker-io / worker-film	job queues by resource: planning and copy; uploads and platform calls; film plans and audits
scheduler	slot decisions, autopilots, keep-ready buffers, generation per day
fleet	starts and stops GPU boxes, assigns jobs to slots, sets boxes up, installs code
sentinel	reaps dead jobs, prunes, alerts, enforces spending caps, runs privacy purges
gateway	a Node.js service, the only door to the models: it speaks the Anthropic Messages API, holds the one Bedrock credential, meters every call per workspace and job, and enforces budgets before a call
voice	the narration service (Chirp 3 HD) the GPU jobs call through a tunnel
mailer	transactional email through Amazon SES, deduplicated alerts
billing	Stripe events and PayPal resync, plan credits, invoices, disputes

Every unit of work is a job with a kind, a resource, a state, attempts and an event log the customer can watch live. 2,735 jobs have run so far, of 47 kinds, from niche suggestions and channel drafts to episode generation, translations, films, insights and comment replies.

Data

PostgreSQL 16 holds workspaces, users, teams, channels, episodes, posts, jobs, credits, plans, invoices, insights and the audit log — 313 tables. It runs on Google Cloud SQL, reached through the Cloud SQL Auth Proxy, with nightly backups and seven days of point-in-time recovery. Redis carries caches, budgets, rate limits and live progress. Media lives in a private Google Cloud Storage bucket, under one prefix per workspace.

CHAPTER 15 · CHANNELFORGE

Multi-tenancy, security and isolation

A SaaS that runs customers' code-generating agents on shared GPUs must keep tenants apart at every layer. ChannelForge does this by design, not by convention.

- **Workspaces, projects, teams and roles.** Every row carries its workspace; teams restrict which projects a member sees; API keys can only be made by owners and admins and carry scopes.
- **GPU job sandboxes.** Each job runs as its own slot user in a private folder no other job can enter. Its network is loopback only — its own tunnels to the model gateway and the voice service; no DNS, no internet — and no job can reach the cloud's metadata service or private networks. A film runs alone on its box, behind a default-deny rule.
- **No secrets on GPU boxes.** The boxes hold no tenant secret and no model key: every model call goes through the gateway, which holds the one real credential and meters and caps each job.
- **Safe outbound calls.** Addresses a tenant chooses — webhooks, Bluesky hosts — are reached only over HTTPS, only at public addresses, and redirects are never followed (no server-side request forgery).
- **Privacy.** A deleted workspace's platform tokens and data are purged 144 hours after the deletion; the YouTube data of a disconnected or refused account is purged on YouTube's own schedule.
- **Abuse controls.** Trial credits are not given to disposable e-mail addresses, and trials are capped per network.
- **Posting guard.** After a database restore the guard arms itself and every account is read before its next post, so nothing posts twice.
- **Accounts.** Every new password is checked offline against 100,000 known-leaked ones; a workspace can require two-factor sign-in for all its members; staff capabilities are checked by every admin action.
- **Data at rest.** The key that seals tenants' platform tokens rotates, with background re-sealing; fetches of addresses a member gives are pinned against DNS rebinding; a member can download all their data, checked against the whole schema.
- **Proven, not assumed.** A planted-defects harness hides deliberate bugs in a copy of the code and checks the tests catch them — 13 of 13 on 8 November — and a guard refuses destructive database commands on production.
- **Operator controls.** An admin console — 42 pages — shows the GPU fleet, model spend, jobs, incidents, publishing failures, finance and the append-only audit log.

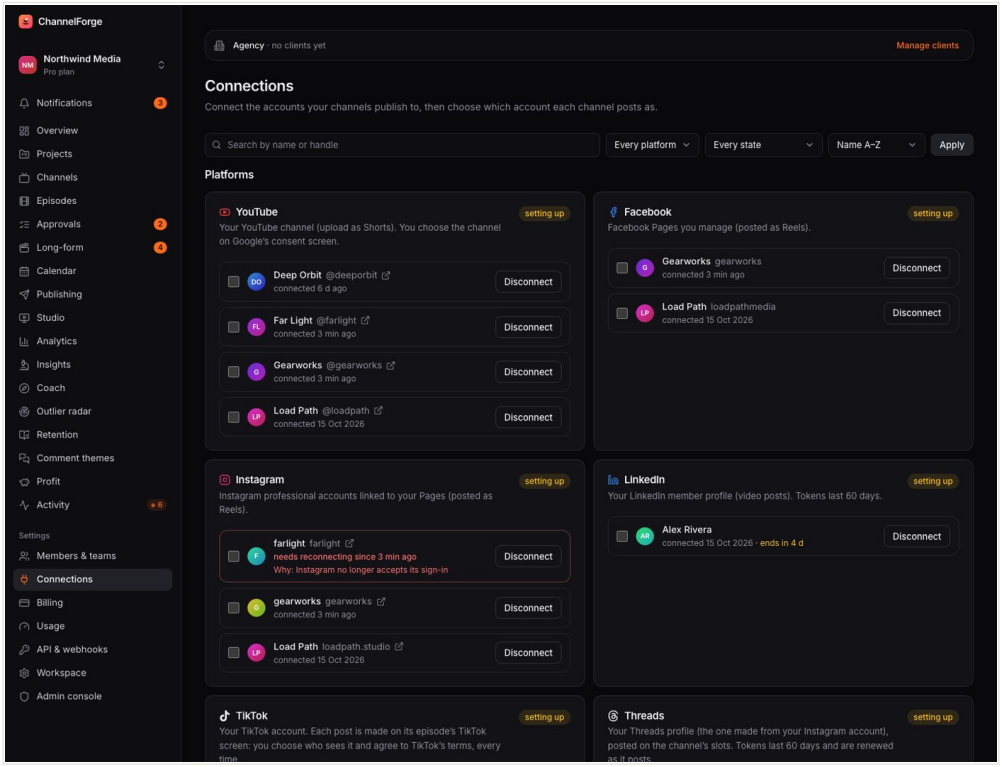


Figure 30. Connections: a workspace links its own platform accounts, each with its health (development build, demonstration data).

CHAPTER 16 · CHANNELFORGE

GPU allocation

GPU time is the scarce resource. ChannelForge treats each L4 box as a set of slots and hands them out fairly.

- **Slots.** Each L4 box offers two slots; a short takes one, a film takes the whole box for its duration.
- **Scaling.** When queued work exceeds free and arriving slots, the fleet starts the cheapest stopped box — cheapest by expected cost per hour, including the risk of pre-emption — with an owner-set preference to break ties or override.
- **Fair share.** A free slot goes to paying tiers first, then to the workspace that used the least GPU time in the last 24 hours for its plan's priority, then to the oldest job — so a big workspace cannot starve a small one.
- **Films and shorts.** A film earmarks the cheapest box and never pre-empts a running job; while shorts wait, films hold at most 60% of the film-capable boxes, because shorts feed posting slots.
- **Blender.** A Blender episode leases a whole box that is set up for Blender, within a share of those boxes and counted against the plan's GPU concurrency, for up to ten hours.
- **Idle stop.** A box with no work is stopped after a short grace period; nothing runs that is not paid for.
- **Failure handling.** A lost box or a refused start is a capacity error: the job is retried elsewhere, never charged.
- **Quality holds.** A rendered episode that fails QA is held and not charged; budgets stop a runaway job with a clear budget error rather than a retry storm.

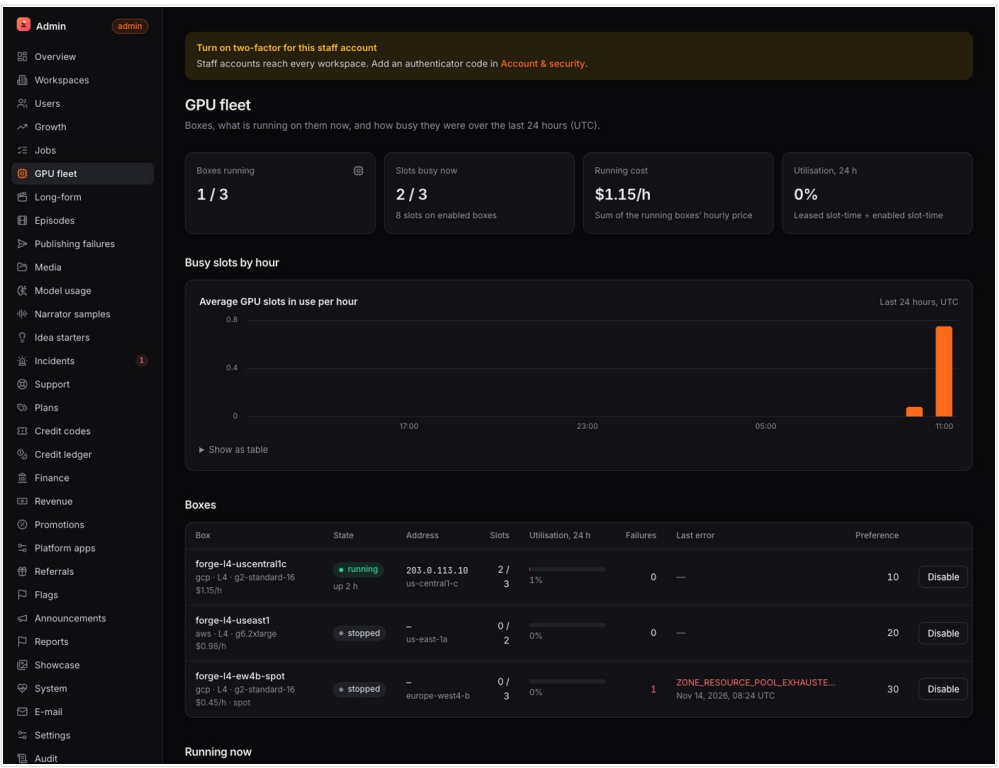


Figure 31. The admin console's GPU page: boxes, slots, utilisation and costs (development build, demonstration data).

CHAPTER 17 · CHANNELFORGE

Plans, credits and billing

Customers buy credits. Every job reserves an estimate when it is queued and is charged what it really cost when it finishes — model calls and GPU seconds at cost, times the plan's markup. A held episode, or a failure on ChannelForge's side, is never charged. Plans add monthly credits and unlock limits and features; packs top up.

Plan	Monthly	Yearly	Credits / month	Channels	Seats	Highlights
Trial	free	–	400 (trial)	2	1	the basics
Creator	$49	$490	600	6	2	Studio, long-form
Pro	$149	$1,490	2,000	24	6	API, Studio, long-form
Agency	$449	$4,490	6,500	100	20	API, Studio, long-form, white label
Enterprise	custom	custom	custom	unlimited	unlimited	API, Studio, long-form, white label

Credit pack	Price	Per 100 credits
500 credits	$55	$11.00
1,500 credits	$150	$10.00
5,000 credits	$450	$9.00
15,000 credits	$1,200	$8.00

One credit is $0.10 of list value. A standard short episode reserves 231 credits when it is queued and is charged what it really cost when it ends. Markups by plan: Trial 3×, Creator 2.5×, Pro 2.3×, Agency 2.1×, Enterprise 1.8×.

Billing machinery

Stripe (subscriptions, checkout, a customer portal of its own, upgrades charged pro rata, downgrades at the period end) and PayPal (subscriptions, refunds and disputes) are built in. Stripe's events are reconciled by polling its Events API, and a returning customer's checkout is applied at once, so a lost webhook can never lose a payment. Selling is behind a switch the owner turns on at launch; until then the free trial is open and the plan and pack catalogue waits in Stripe.

Since the first edition the plans and prices are unchanged, and the money tools around them have grown: credit codes, a time-boxed trial of a paid plan's tools, low-credit and expiry notices, a monthly credit budget per member, and for the operator finance and revenue pages, unit economics, the margin of every job kind and a worst-case break-even per plan. Receipts, promotions, local-currency prices, add-ons and automatic top-ups are built and wait for the launch of selling.

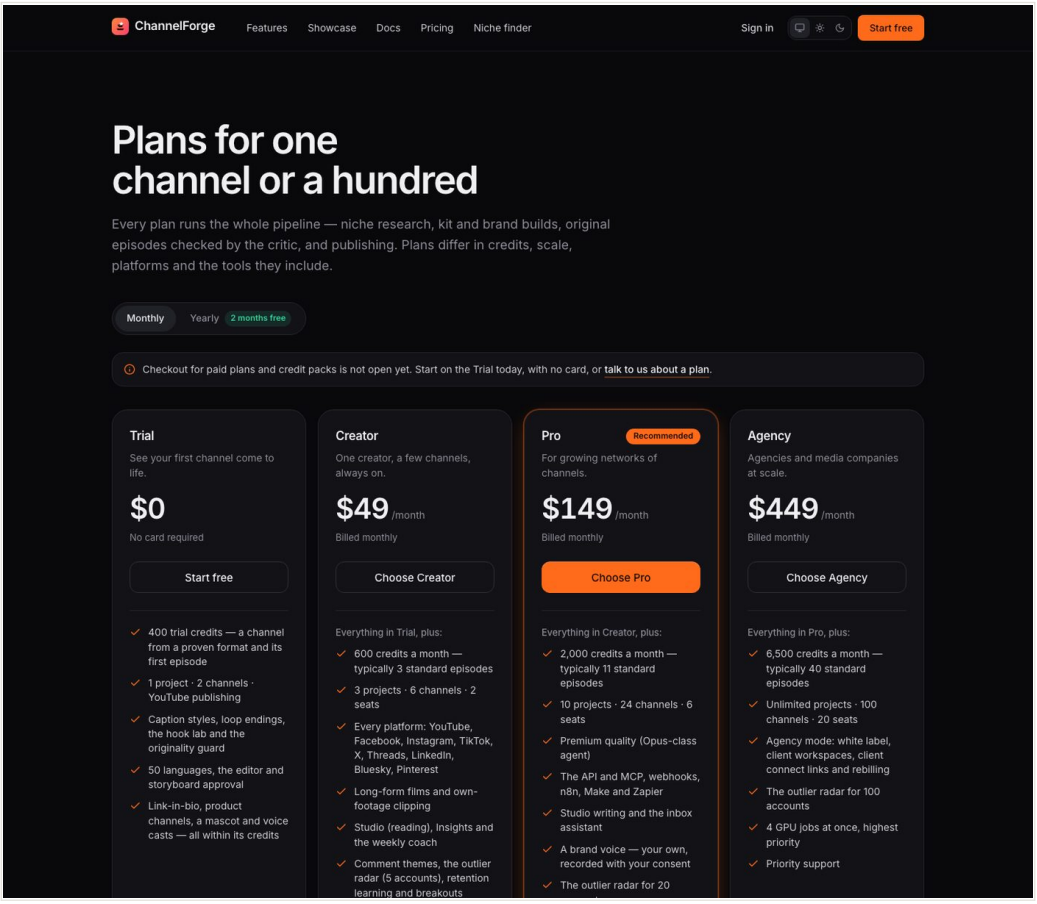


Figure 32. The public pricing page, captured live on 14 November 2026.

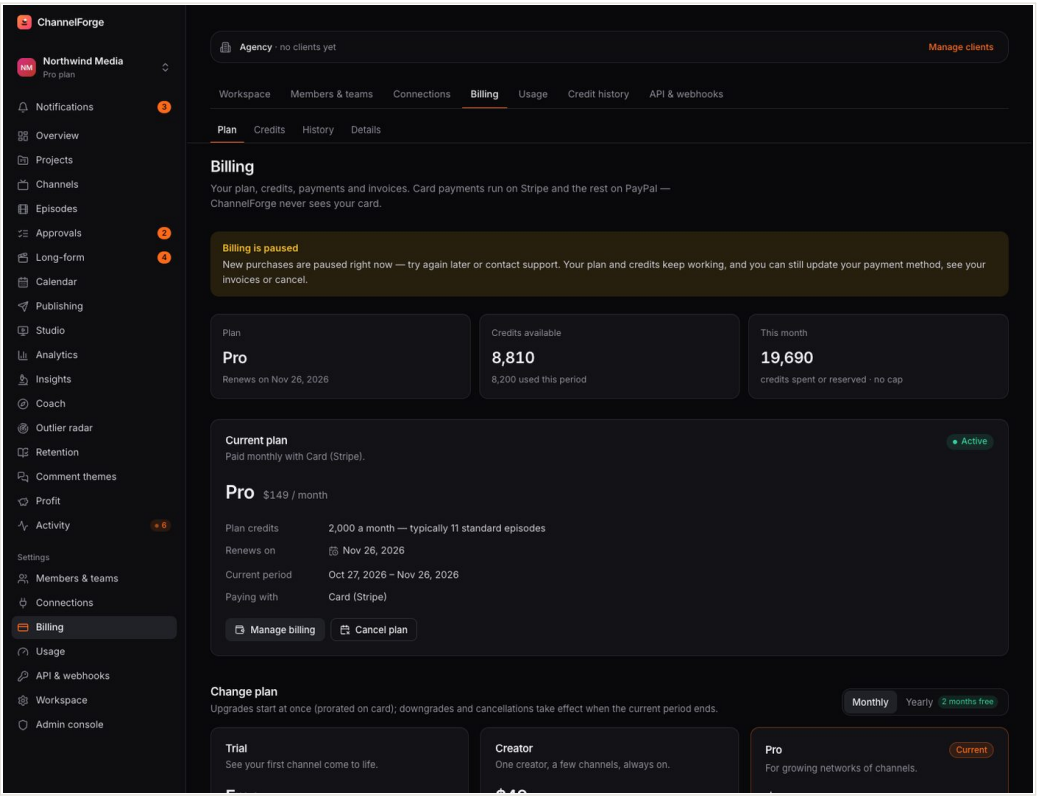


Figure 33. A workspace's billing: its plan, credits and plan changes. 'Billing is paused' is the switch that keeps selling off until launch (development build, demonstration data).

CHAPTER 18 · CHANNELFORGE

API, webhooks and agents

Everything a member can do in the app, a program can do through the API — which is how agencies, automation tools and AI agents put ChannelForge to work.

- **REST API** (OpenAPI 3, 41 route files): channels, projects, episodes, films, jobs, quotes, insights, events, webhooks, style, topic imports, finished-video imports, publishing health; idempotency keys; a quote endpoint that says what a job will cost before it is queued — a Blender episode included; a 14-day request log per API key.
- **Signed webhooks** for job, episode, post and billing events, retried with backoff, delivered without redirects.
- **MCP server**: 30 tools an AI agent (Claude, for example) can call — 'generate episodes for this channel, but spend no more than 300 credits' — each respecting the member's project and channel access; the in-app 'Ask ChannelForge' assistant works through the same tools.
- **Integrations**: n8n, Make and Zapier connectors speak the same API.

Example: queue one episode with a credit ceiling

```
POST /api/v1/episodes/generate
Authorization: Bearer <api key>
{"channel_id": "ch_…", "max_credits": 231, "count": 1}
→ 202 {"queued": [{"job_id": "job_…", "credits_held": 231}], …}
```

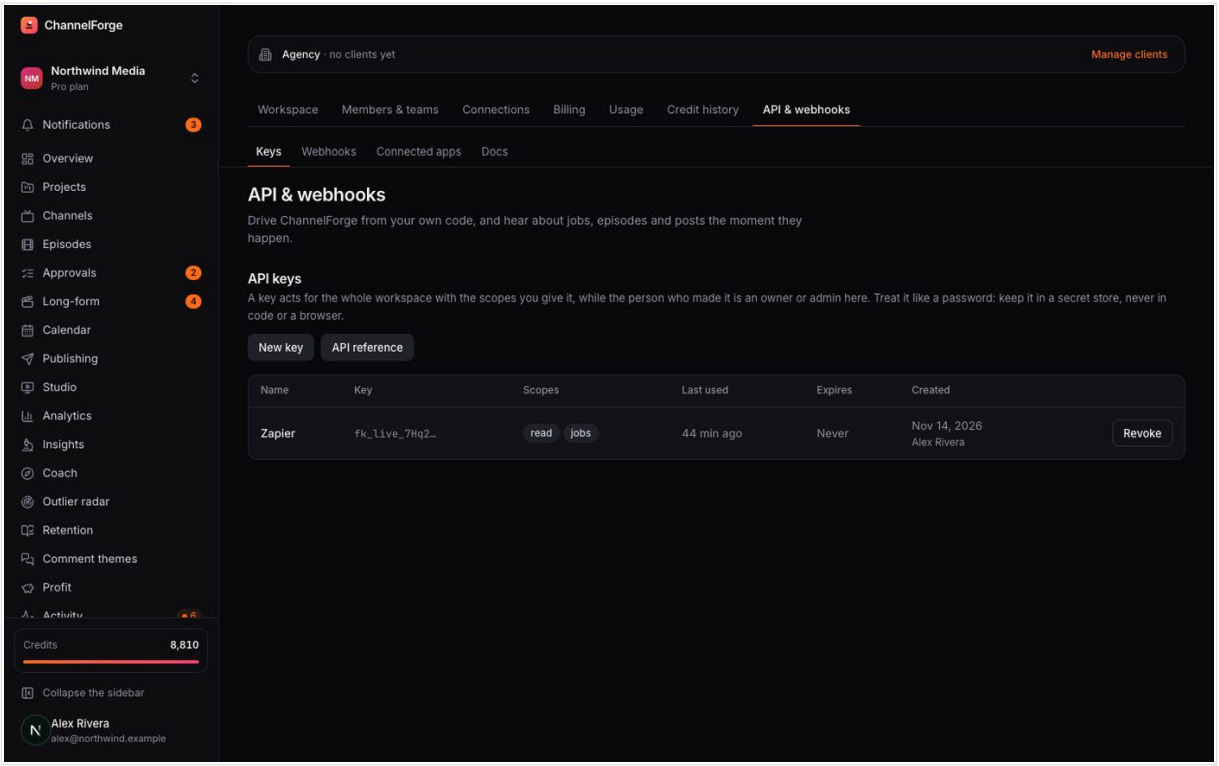


Figure 34. API keys and webhooks in a workspace (development build, demonstration data).

PART IV

Infrastructure, operations and economics

Where everything runs, how the GPUs are managed, how nothing is ever lost, what it is built with, what it costs, and how long it took.

CHAPTER 19 · INFRASTRUCTURE

Infrastructure map

Both systems run on Google Cloud in europe-west1 since 2 November 2026, with their GPUs on AWS in us-east-1 and us-east-2 and their models on Amazon Bedrock. Cloudflare fronts both sites.

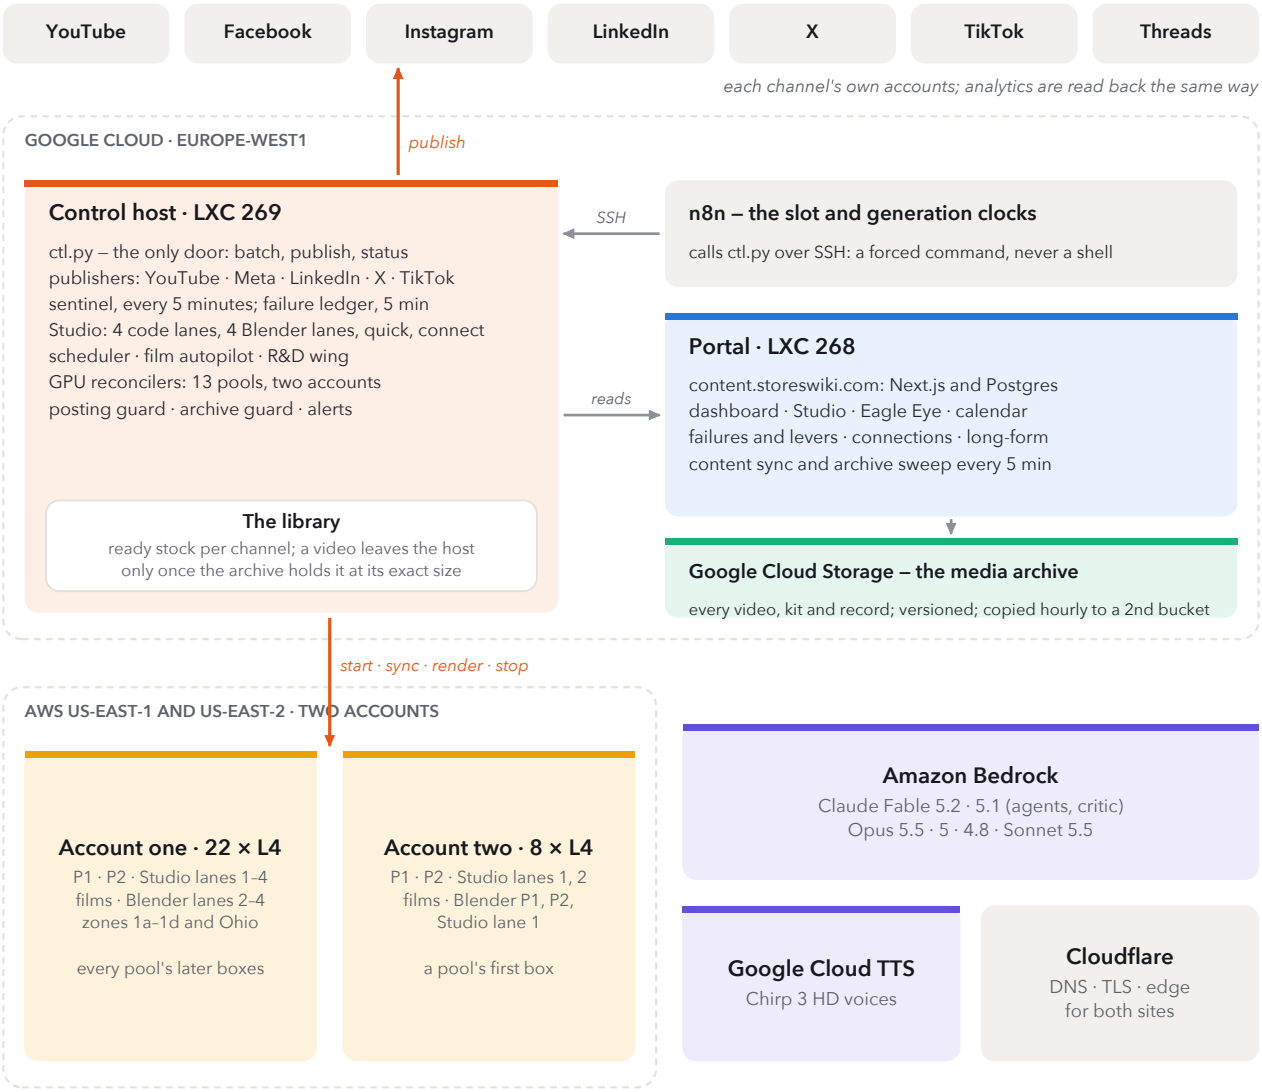


Figure 35. The original programme's architecture. ChannelForge (Figure 29) runs on its own host, database and GPU fleet.

Component	Where	What
Control host (original)	Google Cloud, a Linux container on the estate's Proxmox host	ctl, publishers, sentinel, schedulers, workers, reconcilers, the failure ledger
Portal (original)	Google Cloud, a second container beside it	content.storeswiki.com, its Postgres, content sync
ChannelForge host	Google Cloud, an Arm VM (c4a)	app, control plane, gateway, Redis
ChannelForge database	Google Cloud SQL, PostgreSQL 16	313 tables, through the Cloud SQL Auth Proxy
Media	Google Cloud Storage, a bucket per system	the original's versioned, Autoclass, 90-day soft delete, copied every hour to a second bucket; ChannelForge's private, per workspace
GPUs	AWS us-east-1 and us-east-2, two accounts	NVIDIA L4: 30 boxes for the original in 13 pools, ChannelForge's own fleet beside them
Models	Amazon Bedrock	Claude Fable 5.2 and 5.1, Opus 5.5, 5 and 4.8, Sonnet 5.5, Haiku 5.5
Voice	Google Cloud Text-to-Speech	Chirp 3 HD; word timing by faster-whisper on the GPUs
Schedules	n8n	slot and generation clocks for Projects 1 and 2
Edge	Cloudflare	DNS, TLS, proxying for both sites
Email	Amazon SES	alerts and ChannelForge's mail

CHAPTER 20 · INFRASTRUCTURE

The GPU fleet

Rendering and the episode agents run on NVIDIA L4 GPUs (24 GB, Ada generation) in g6 machines on AWS. Each is a software twin made from a golden image — the same driver, Chrome, Blender, Node, Python packages, speech recogniser and agent CLI on every box — so a scene renders identically anywhere. The original programme's fleet has grown from 12 boxes on 3 November to 30, in 13 pools: one each for Projects 1 and 2, four Studio lanes, the films, and six for Blender episodes.

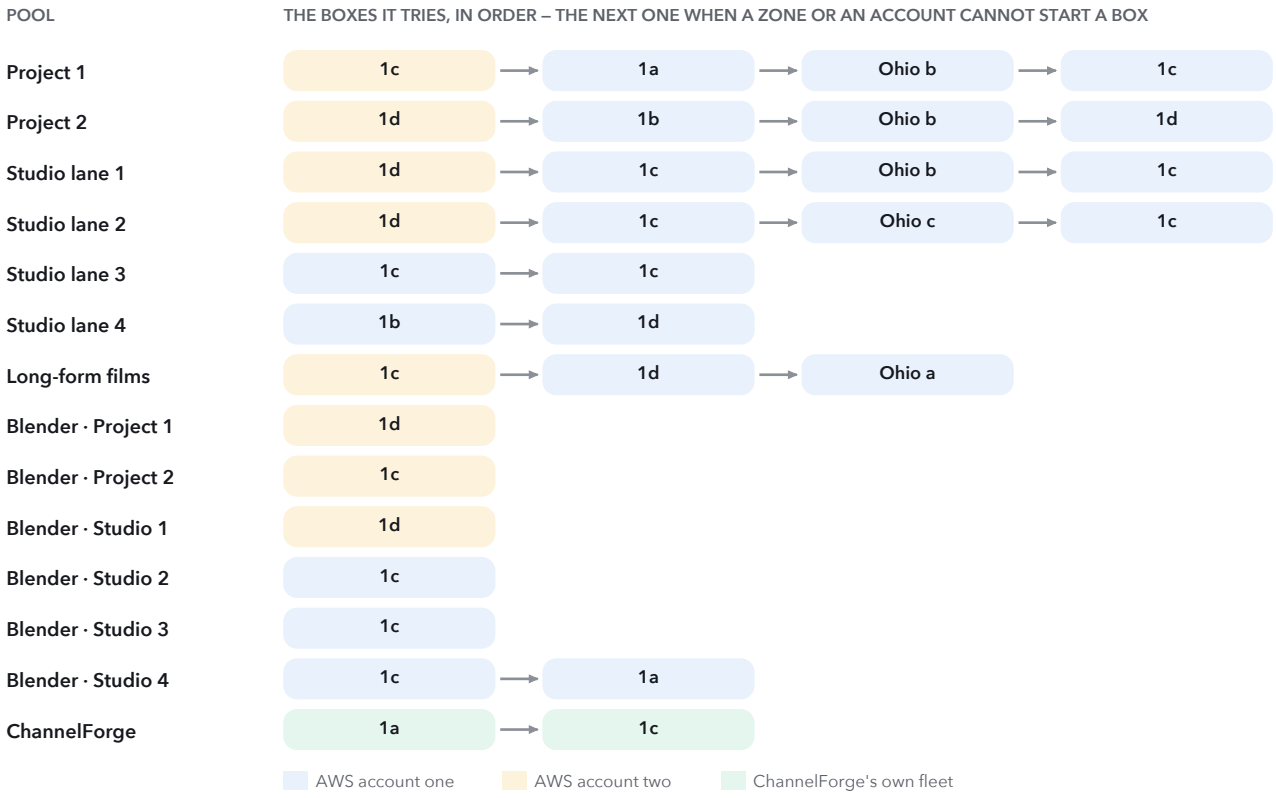


Figure 36. The original programme's 13 GPU pools and ChannelForge's fleet: the boxes each pool tries, in order, by zone (1a–1d: us-east-1's zones; Ohio: us-east-2) and by account. A zone without L4 capacity, or an account that refuses the start, sends it to the pool's next box.

Reconcilers: GPUs on only while working

Work declares intent; a reconciler per pool turns intent into machines. A batch or job keeps touching an intent file while it needs a GPU; every minute the pool's reconciler makes the provider state match: a fresh intent starts a box (or keeps one running), a stale one stops it. Nothing has to remember to clean up — silence means off — and absolute caps stop any box that runs too long. A box serving a batch is never switched away from mid-run, and a box's SSH host keys are pinned by name so a recycled address can never reach the wrong machine.

Capacity and quotas

Cloud GPUs are scarce: on 3 November alone, L4 starts were refused in three us-east-1 zones at different times of the day. The fleet is built for it: boxes sit in different zones, regions and accounts, a refusal moves on to the pool's next box, and a refusal of the account itself is treated the same way. The work is fitted to the quotas there are rather than the other way round: the Studio's newest projects use lanes of their own, kept off the live projects' boxes, and give way to them whenever a live batch has waited eight minutes. Every new box is checked — GPU, driver, Vulkan, software versions — and proven by rendering reference stills byte-identical to the fleet's references before it joins a pool.

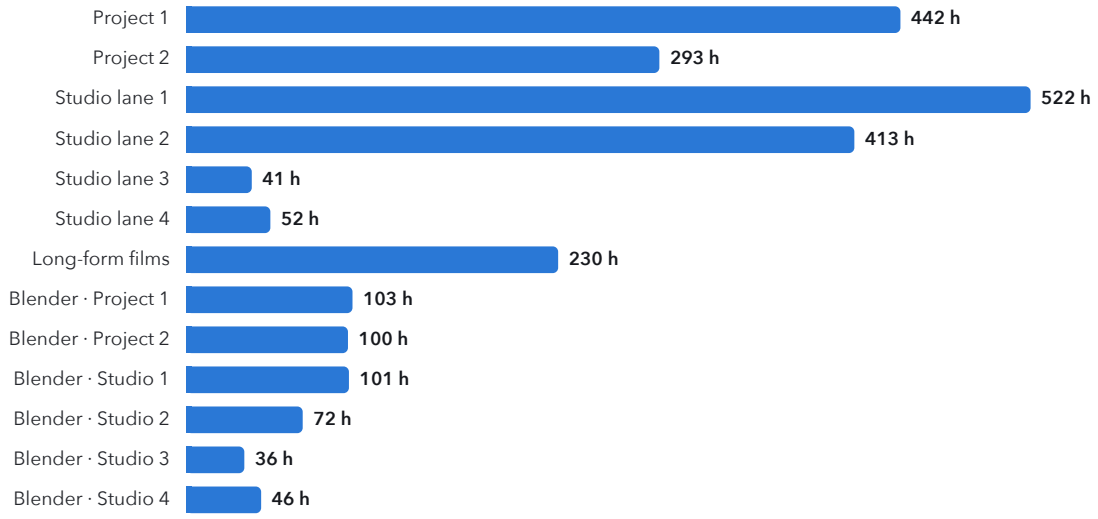


Figure 37. Billed L4 hours by pool, 12 October – 14 November 2026: 2,451 hours in all. The Studio's first lane has run the most; the Blender pools began on 10 November.

CHAPTER 21 · INFRASTRUCTURE

Media, backups and disaster recovery

The owner's rule: everything made is kept, in more than one place, and the whole system can be rebuilt from its backups on new infrastructure — turn-key.

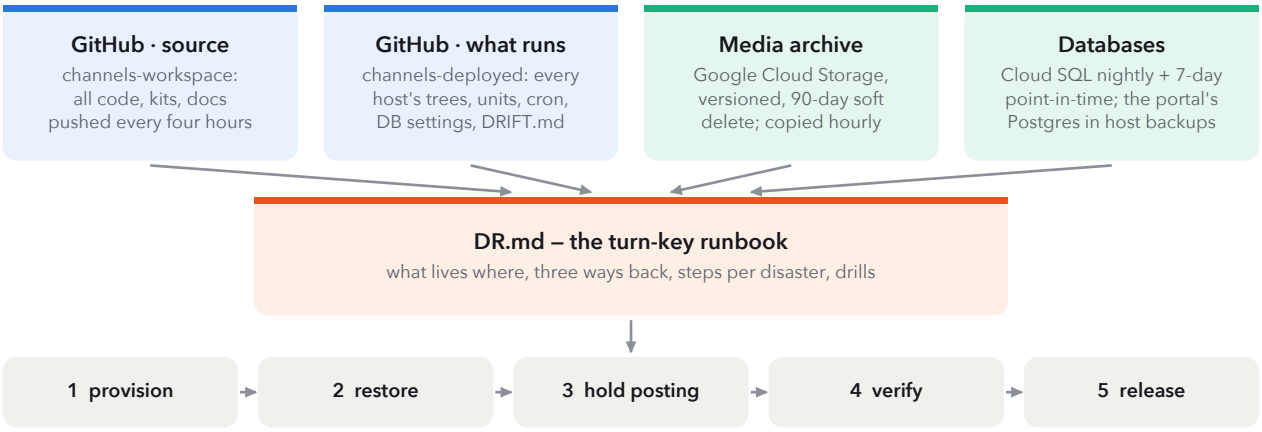


Figure 38. What the systems are recovered from, and the order a recovery follows.

The media archive

Every video since the first day is in Google Cloud Storage: all 2,015 episodes of the original programme with their Facebook edits and thumbnails, every film, every kit, Blender asset and brand source, every record — 730 GB in all — and all 944 of ChannelForge's episodes in its own bucket. The original's archive keeps every overwritten version, and anything deleted from it can be restored for 90 days. Every hour a transfer job copies both buckets, whole, into a second pair of buckets kept in the archive storage class, so every video exists twice. The control host deletes a local video only when the archive holds it at its exact size.

Code and configuration

Two private GitHub repositories are refreshed every four hours and after every deploy (every night until 8 November): the workspace (every line of code, the kits, the brand art, the documentation) and a snapshot of what actually runs on every host — code trees, services, timers, cron, database settings, GPU inventory, the workflows — with a drift report against the workspace. A secret scanner, which knows the estate's real secrets and the shapes of every kind of key, stands in front of both pushes; nothing secret is ever committed. Since 14 November both repositories are also kept as git bundles in a cloud bucket, so even a lost GitHub account would lose nothing.

Golden images in both clouds

Every GPU box is made from a golden image, and the images are kept on both clouds: on AWS in two accounts, and on Google Cloud, where the image was rebuilt on 14 November with the whole current stack — the driver, Chrome, Node, the agent CLI, Blender, the word aligner and its weights, the 2.9 GB library of CC0 assets — and checked against a live AWS box: the same versions, and seven test renders identical to the pixel. What a new box would otherwise fetch from third parties (the Blender release, the asset library, the aligner's weights) is stored with its checksum in a bucket on each cloud. If one provider refused every start, both systems' boxes could be made on the other from the same recipe, in minutes.

The runbook

DR.md describes where everything lives, the three ways back (standby, backups, rebuild from GitHub), the steps for each kind of disaster — a lost GPU box, a lost host, a lost database, a lost cloud account — and the drills that prove them. A restored host starts with posting held; the posting guard reads every platform before the first post; only then is posting released.

CHAPTER 22 · STACK

Technology stack

Chosen for reliability and for what each does best; nothing exotic, everything replaceable.

Layer	Technology
Languages	Python 3 (control planes, pipeline, tools), TypeScript (web apps, the model gateway), JavaScript (renderers, kits), SQL, Bash
Web	Next.js 15, React 19, Tailwind CSS 4, Prisma, nginx, Cloudflare
Data	PostgreSQL 16 (Cloud SQL), Redis, JSON ledgers and records on disk, Google Cloud Storage (S3 API)
AI models	Claude Fable 5.2 and 5.1 (agents, critic, films), Claude Opus 5.5, 5 and 4.8 behind them in a fallback chain, Claude Sonnet 5.5 (ChannelForge's standard agent, copy, planning), Claude Haiku 5.5 – on Amazon Bedrock, across two cloud accounts
Agent runtime	Claude Code CLI in per-job sandboxes on the GPU boxes
Rendering	headless Google Chrome, WebGL2 and Canvas, NVIDIA L4 via Vulkan/ANGLE, ffmpeg; Blender 4.5 LTS (EEVEE, headless through EGL) with Poly Haven CC0 assets and OpenStreetMap data
Voice and audio	Google Chirp 3 HD voices, faster-whisper large-v3 for word timing, Google speech-to-text for clipping, synthesized music and effects, loudness to −14 LUFS
Orchestration	systemd services and timers, n8n workflows, job queues in Postgres and on disk
Cloud	Google Cloud (Compute, Cloud SQL, Storage, TTS), AWS (EC2 GPUs, Bedrock, SES), Cloudflare
Payments	Stripe, PayPal
Platforms	YouTube Data API, Meta Graph (Facebook, Instagram, Threads), LinkedIn, X API v2, TikTok, Bluesky, Pinterest, Mastodon, Telegram
Quality	unit and integration tests (1,106 test files), rendering parity checks, browser checks of every page and form, a 'nothing lost' suite for every page's tabs

Size of the code

Code tree	Lines of code	On 3 Nov	Files	Docs (lines)	Test files
Original pipeline, control plane and portal	264,396	212,000	1,079	25,220	183
Studio backend (Projects 3–10)	172,134	142,000	668	16,081	83
ChannelForge (app, control plane, gateway, pipeline)	708,460	346,000	3,981	83,636	840
All three	1,144,990	700,000	5,728	124,937	1,106

Counted on 14 November 2026, the same way as for the first edition: without dependencies, builds or media. The Studio backend began as a copy of the original pipeline and ChannelForge's pipeline as a copy of the Studio's, so the trees share a common core.

CHAPTER 23 · ECONOMICS

Costs and unit economics

Three things cost money: model calls, GPU hours and the hosting around them. The first two scale with output; the third is small and nearly fixed.

Per episode (measured)

	Original programme	ChannelForge
Model time per short episode	$10.82 on average (Project 1 $10.07, Project 2 $9.99, Studio $11.88); $9.57 an episode on 3–9 November, $15.52 on 10–13 November, on Fable 5.2	$7.02 (standard quality: a Sonnet 5.5 agent)
A Blender episode	$11.39 in model time, plus 1.5–5 L4-hours of render	a price class of its own: about 700–1,100 credits at Pro
GPU time per episode	about 1.2 L4-hours all in (boot, set-up, Blender and films included): $1–2	$0.56 on average, at cost
Length of a short	142 s on average (plus a ≤ 89 s edit)	148 s on average
A long-form film	about $156 in model time on average, several GPU hours	priced per film from a quote

To date (measured)

	Original programme	ChannelForge
Model calls	$21,805 (short episodes, Blender included)	$9,891 (all job kinds)
GPU hours	2,451 h across the pools	1,240 h ($529 at cost)
Output	2,015 episodes, 28 films	944 episodes, 3 films

Running costs (estimates)

Item	Estimate	Notes
Hosting: two containers, ChannelForge's VM, Cloud SQL	$150–300 / month	small machines; most of the work is on GPUs
Media storage	$15–40 / month	730 GB on 14 November 2026, against 308 GB on 3 November; Autoclass moves what nobody reads to colder classes, and the hourly copies sit in the archive class
GPU boxes while stopped	about $10 / month each	the 120 GB disk; 30 boxes ≈ $300 / month
GPU hours while working	$0.98–1.32 / hour per L4	on-demand; a box works only while episodes wait
Models	$10–18 per short episode	the main variable cost; budgets cap each project's day
Voice, e-mail, DNS	a few dollars / month	Chirp per character, SES per e-mail

Where the money went

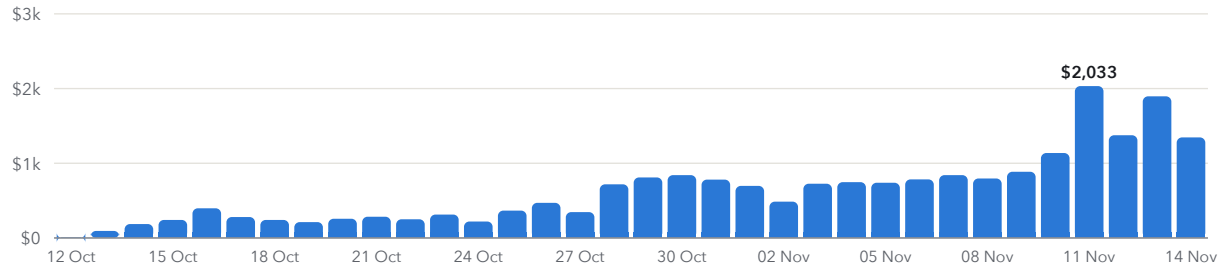


Figure 39. Model spend per day of the original programme, in US dollars (12 October – 14 November 2026; the last day is partial). The step on 10–11 November is the move to Claude Fable 5.2 and the first Blender episodes.

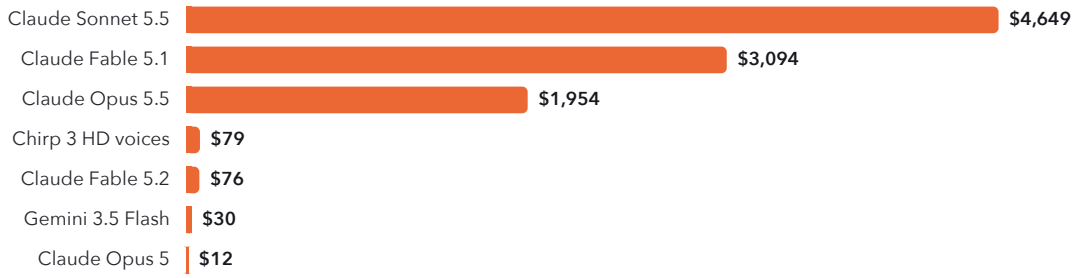


Figure 40. ChannelForge's model and voice spend by model, all job kinds, 18 October – 14 November 2026 ($9,895 in all). Standard episodes run on Sonnet 5.5; the owner's workspace has used Fable 5.1 agents and Fable 5.2 reviews since 10 November; premium ones and films run on Opus 5.5.

What one short earns on ChannelForge

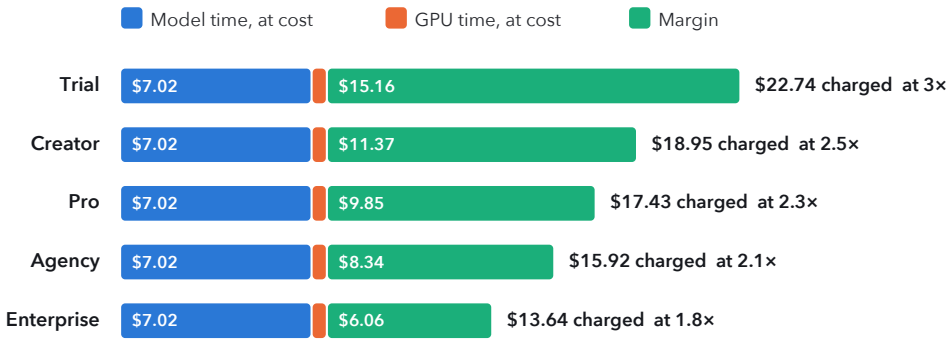


Figure 41. One standard short on ChannelForge: what it costs — $7.02 of model time and $0.56 of GPU time on average, both at cost — and what each plan charges for it: the cost times the plan's markup. The margin is earned on every job, before any subscription.

UNIT ECONOMICS

In the four days 10–13 November the original programme made 79 to 129 episodes a day for $1,137–$2,033 a day in model time, plus one to two dollars of GPU time an episode — for 166 to 226 platform posts a day, the mirrors and a film.

ChannelForge charges the same work at cost times 1.8 to 3. At list value that is a margin of 44–67% on every credit; at the plans' own credit prices a fully used plan still earns 51% (Creator), 42% (Pro) and 31% (Agency), and packs 40–64%.

CHAPTER 24 · HISTORY

Timeline and effort

From nothing to two production systems in 33 days: designed and directed by the author, and built with AI engineering agents (Anthropic's Claude) working to that design.

33 days
from the first episode to this edition

1.14M
lines of code across the three trees

1,106
test files

1,527
commits since the first backup

Date (2026)	Milestone
12 Oct	The first episodes: 'What if Yellowstone erupted 500 miles from you?' and a Category 6 hurricane
13 Oct	Project 1 live: four channels on YouTube, Facebook and Instagram, slots on n8n
14–15 Oct	Project 2's first episodes; its twelve accounts live on 15 October; the self-healing sentinel
15–17 Oct	The Studio: Projects 3 and 4 created through the browser, and their first episodes
18 Oct	ChannelForge live: multi-tenant app, control plane, GPU fleet, billing; nightly GitHub backups
19 Oct	Long-form films; GPU placement across pools
20–21 Oct	ChannelForge's browser end-to-end test (16 channels, 45 episodes, 2 films); video format v4 with Chirp 3 HD voices; X
22 Oct	Project 4's first four channels live; the LinkedIn queue; TikTok and Threads apps
25–26 Oct	The film autopilot (one a day); Projects 5 and 6 (sixteen new channels); Model Cutaway, the 13th live channel; the inbox assistant
27–29 Oct	Self-rescue systems: failure e-mails, a film doctor, box-loss recovery; every video kept in two clouds
31 Oct	ChannelForge market research, and 25 new features built in three waves
1 Nov	Everything switched on; turn-key disaster recovery; ChannelForge's public site and 75 docs pages; the posting guard
2 Nov	The whole estate moved to Google Cloud: hosts, database and media
3 Nov	GPUs on AWS only, in two accounts with zone failover; a full audit and rehearsal; the first edition of this dossier
4 Nov	Covers on every platform; the portal's calendar and publishing pages from 12–18 s to under half a second; a bigger L4 in every pool
5 Nov	Project 7 (eight channels) built in the Studio; a GPU region in Ohio for every pool; ChannelForge's first feature study and its first wave live
6 Nov	Live episodes paired two to a GPU box; model calls fall back across two cloud accounts
7 Nov	Held films found and the film ladder fixed; ChannelForge's second study, 85 features, live
8 Nov	Every video in two buckets, the GitHub backup every four hours; X's 140-second rule; ChannelForge's pages in tabs and its third study live
9 Nov	The YouTube processing watch; the Blender research film, rebuilt from real assets and word-timed captions
10 Nov	Blender episodes live in both systems, the first Blender slots the same day; agents on Claude Fable 5.2
11 Nov	Every live channel posts Blender episodes twice a day; Blender agents on Opus 5.5
12 Nov	Projects 8, 9 and 10 (24 channels) and three new ChannelForge projects (24 more); films for channels not yet live; the 'nothing toy-like' rule for Blender
13 Nov	Captions that are the voice, live for every project's shorts; failures with levers in both systems; the portal's pages in tabs; Connections and the live schedule; old-format videos held at launch; Signal Voyage and Unseen View live
14 Nov	Upright Robots live; GPU fail-over across cloud accounts proven; a fairer release floor for films; this second edition

PART V

Results and the offer

What the channels have reached so far, what a buyer receives, and who built it.

CHAPTER 25 · RESULTS

Audience and results

A month is early for a channel; most take months to find their audience. Still, every live channel has an audience, several have had videos travel far beyond their followers — one Into The Deep video, on the 28,800 bath toys that mapped the Pacific's currents, passed 740,000 views on Facebook alone — and the system has learned which subjects travel.

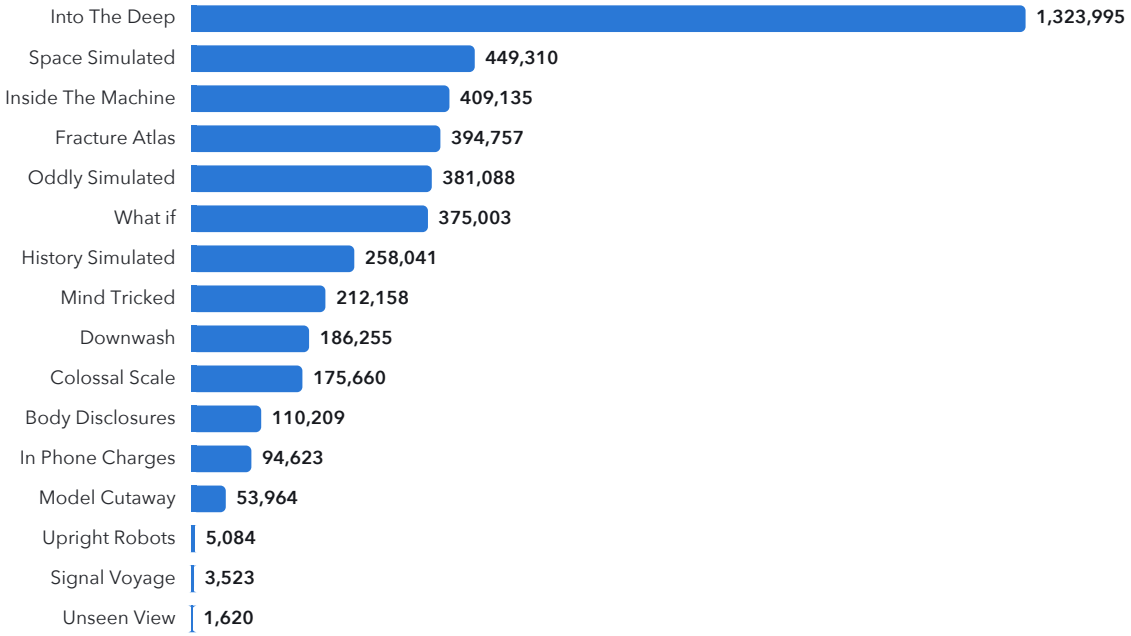
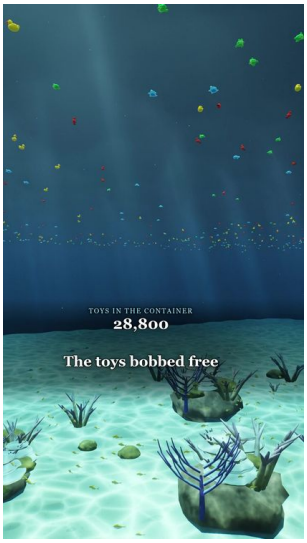
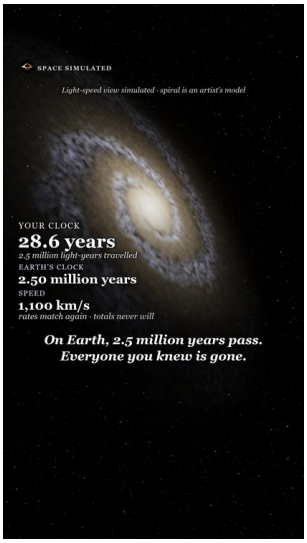


Figure 42. Views per channel, all platforms, 16 live channels (14 November 2026). The three newest went live on 13 and 14 November.

The videos that travelled



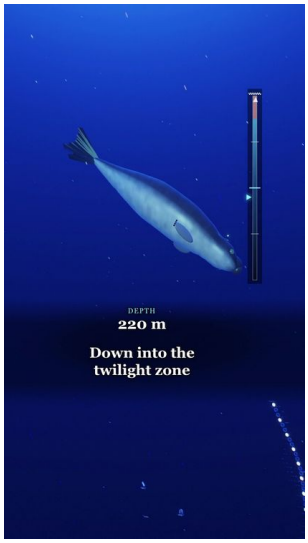
743,468 views · Into The Deep
28,800 bath toys fell into the Pacific. They mapped its currents



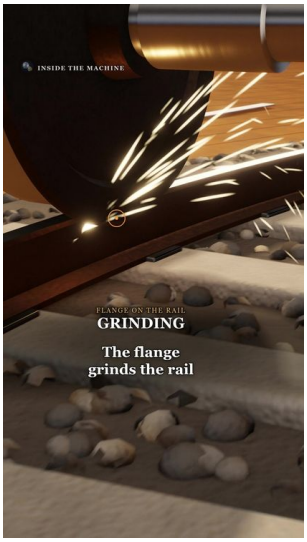
116,260 views · Space Simulated
Could you reach Andromeda in one lifetime?



83,777 views · What if
What if an EF5 tornado came down your road?



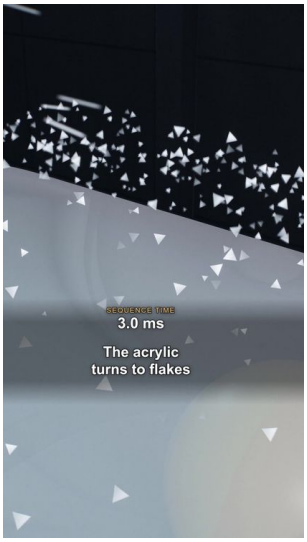
63,541 views · Into The Deep
This seal falls asleep mid-dive and spirals down, upside down



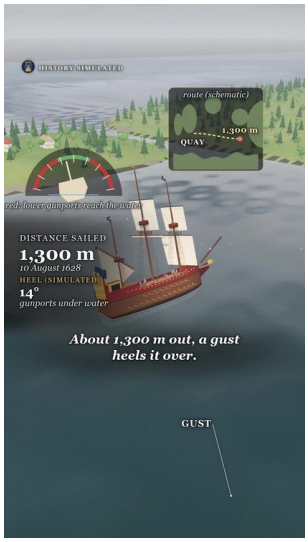
48,526 views · Inside The Machine
Train wheels share one solid axle. So how do they turn?



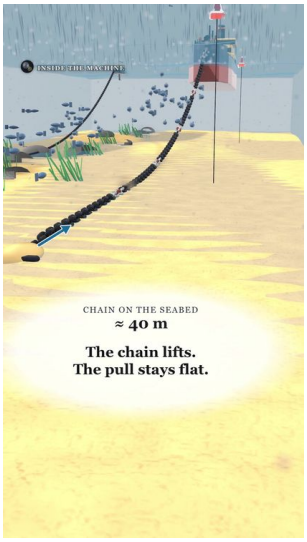
48,349 views · Oddly Simulated
Can a 5 mm domino topple one as tall as a skyscraper?



47,352 views · Fracture Atlas
Why do some jet canopies carry a line built to shatter them?



35,465 views · History Simulated
This warship failed its stability test, then sailed anyway



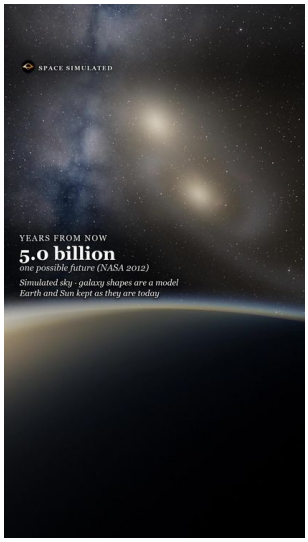
33,962 views · Inside The Machine
Why does a ship let out far more chain than the sea is deep?



29,009 views · Colossal Scale
This 7 km train has engines in the middle, run by radio



28,956 views · Fracture Atlas
Why car springs get blasted with millions of tiny steel balls



26,802 views · Space Simulated
What will the sky look like when Andromeda arrives?

Figure 43. The most-watched videos, with their real thumbnails (views on their best platform).

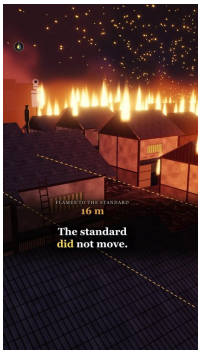
Followers by channel

Channel	YouTube	Facebook	Instagram	Total
Body Disclosures	45	23	66	134
Space Simulated	270	96	523	889
History Simulated	8	44	363	415
What if	552	19	102	673
Inside The Machine	483	139	101	723
Mind Tricked	80	47	21	148
Into The Deep	182	8,850	92	9,124
Oddly Simulated	213	25	69	307
Colossal Scale	56	136	39	231
Downwash	110	19	100	229
Fracture Atlas	145	165	158	468
In Phone Charges	30	10	44	84
Model Cutaway	10	4	17	31
Unseen View	3	0	0	3
Signal Voyage	7	0	2	9
Upright Robots	4	6	3	13
All live channels	2,198	9,583	1,700	13,481

Quality

The average critic score of everything the original programme made is 8.08 out of 10, and only 30 of 2,015 episodes (1.5%) were held by the gates, with 23 more rejected by a review — 22 of them Blender episodes turned down while the format found its look. ChannelForge's episodes average 8.05. Every one of these numbers is enforced, not hoped for: an episode under the bar cannot reach a platform.

A gallery of the channels



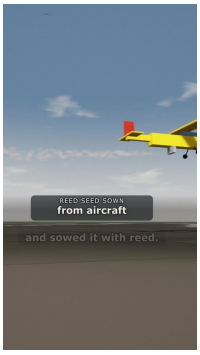
Beforehand
Edo's firefighters stopped fires by tearing…



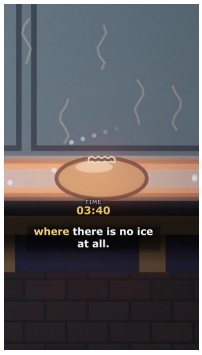
Brake Run
The brake that needs no power and can never…



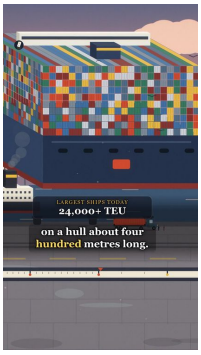
Why That Size
Earth is 40,000 km around because of how we…



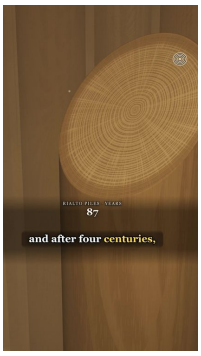
Cofferdam
They drained a sea, then sowed the mud with…



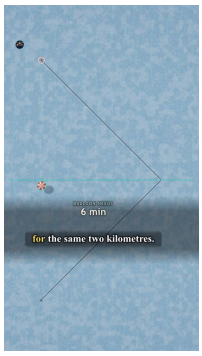
Below The Block
Frozen pipes usually don't burst where the…



Domino Trail
16 cents a ton: the 1956 box that made…



Endgrain
A fence post can rot in 10 years. Venice's…



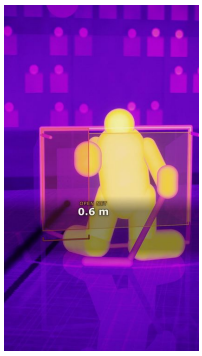
Engineless
Downwind, an iceboat beats the balloon…



Envenom
The mouse that turns scorpion venom into a…



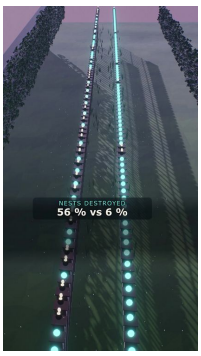
Fledge
Flamingo milk is red. Feeding it drains the…



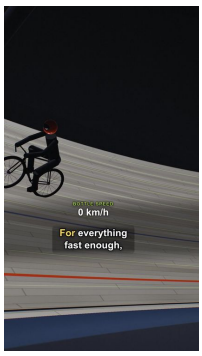
Halfspace
A pass moves the goal faster than a goalie…



Heatwork
Why a hot dish cracks on a wet counter (and…



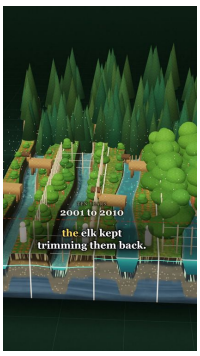
Host Hijack
Throw out the cowbird's egg and the nest…



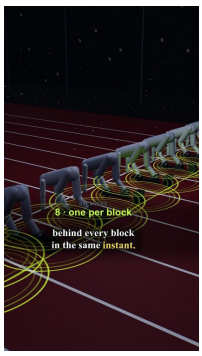
Lean Angle
To a rider at 50 km/h, this 45° wall is a…



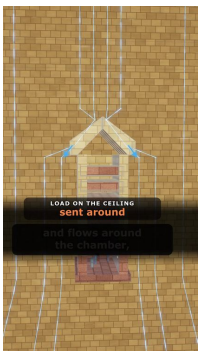
Odd Ledger
Why coins have ridges – and why old houses…



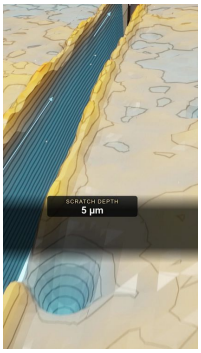
Linchpin
Wolves came back. The Yellowstone willows…



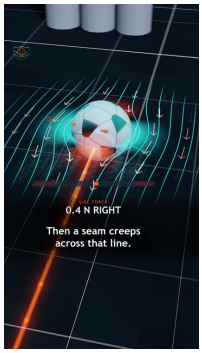
Line Call
The starting gun is a loudspeaker: a bang…



The Load Path
Five empty rooms keep the Great Pyramid…



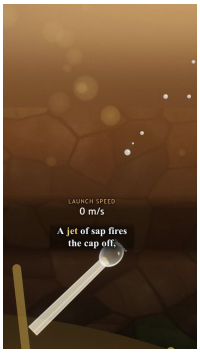
Within a Micron
Your tap drips because of a scratch you…



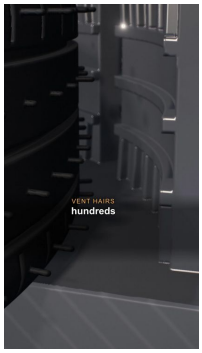
Midflight
Kicked with almost no spin, a football…



Mutual Terms
Each dawn this squid evicts 95% of its…



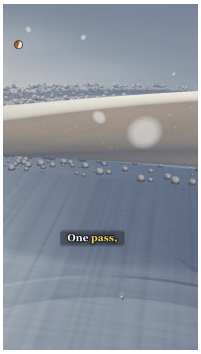
Mycelium Minute
This dung fungus aims at the sunrise and…



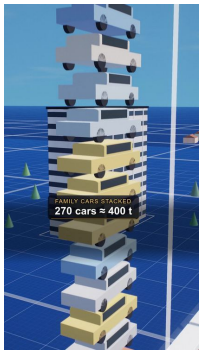
Off The Line
The hairs on a new tyre are the mould's…



Outlived
The whale that swam into 2007 with an 1880s…



Under the Patina
One wipe erases this photo: its face is…



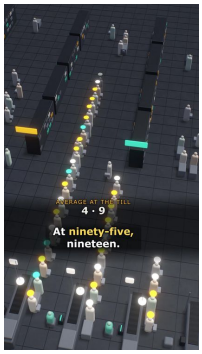
Railhead
A 400-tonne train stands on one playing…



Specimen Bench
A wire this thin lifts a full mug. Its…



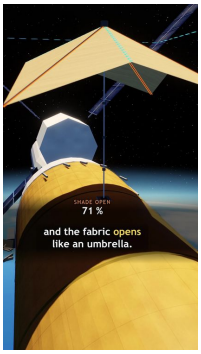
Tenths
In the rain, the fastest racing line…



Throughput Lab
A 5-minute till break, and a 45-minute queue



Unbuilt
The 140 m Tokyo hotel that shrank two…



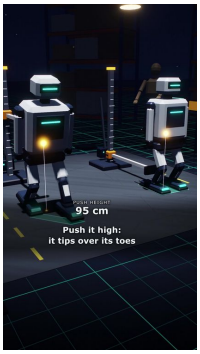
Unfurled
A space station saved by a parasol pushed…



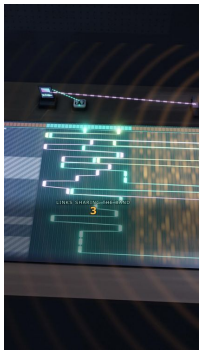
Now You Don't
Zebra stripes aren't for lions. Watch a…



Unseen View
Dip your hand in water and a thermal camera…



Upright Robots
Push a robot low and it slides. Push it…



Signal Voyage
Bluetooth can change channel up to 1,600…



In Phone Charges
A rifle bullet carries the energy of a car…

Figure 44. One recent episode from each of the 32 channels added since the first edition (Projects 7 to 10, not yet live), from the three that went live since — Signal Voyage, Unseen View and Upright Robots — and from In Phone Charges.

CHAPTER 26 · THE OFFER

What a buyer receives

Either system can be acquired on its own, or both together. Three forms are possible: outright purchase with full rights, a licence to run it, or a done-for-you build on the buyer's own cloud accounts with the buyer's own channels.

	Original programme	ChannelForge
Source code	pipeline, kits, control plane, portal, Studio	app, control plane, gateway, pipeline, infrastructure
Running system	65 channels with their brands, kits and stock, the Blender toolkit and its asset catalogue; optionally the live accounts	the live product, its flags, plans and catalogue
Machine images	golden GPU images (both clouds), setup scripts	the GPU box images and setup
Data	the media archive and its replica, analytics history, R&D findings, the failure ledger	the schema, migrations, settings, the three feature studies
Documentation	READMEs, runbooks, DR.md, the operations lessons, this dossier	the same plus 82 public docs pages, an operator RUNBOOK and the three feature studies
Handover	guided transfer of hosts, accounts and schedules	guided deployment and the first weeks of operation

What the buyer needs

- A cloud account for hosting (Google Cloud as today, or any Linux cloud) and one or two AWS accounts with an L4 GPU quota.
- Access to Claude models (Amazon Bedrock, or Anthropic's API) and Google Cloud Text-to-Speech.
- Developer apps on the platforms to publish to: a Google Cloud project for YouTube, a Meta app for Facebook, Instagram and Threads, and LinkedIn, X and TikTok apps as needed — each needs the platform's review for public use.
- For ChannelForge: Stripe and/or PayPal merchant accounts, a domain, Amazon SES (or any SMTP provider).

How it would be moved

1. Provision hosts, a database and buckets on the buyer's cloud from the infrastructure scripts.
2. Restore code from the repositories and the data from the backups; install the golden GPU images.
3. Hold posting; connect the buyer's model, voice and platform accounts; run the account checks and a render parity check.
4. Rehearse: one generation per pool, one publish per platform, a schedule cycle — then release posting.
5. Hand over the runbooks and operate together for an agreed period.

CONTACT

Majid Khan Mohmand – WhatsApp +39 331 106 0553 · linkedin.com/in/majidkhanmohmand · fiverr.com/majidkhan_moh

Via Giuseppe Mazzini 4, 95032 Belpasso (CT), Italia

CHAPTER 27 · THE AUTHOR

About the author and the ecosystem

Majid Khan Mohmand designs and builds AI-driven products and the infrastructure under them, and runs an ecosystem of software-as-a-service products for creators, marketers and small businesses. These video studios are its newest members.

LinkedIn	linkedin.com/in/majidkhanmohmand
Fiverr	fiverr.com/majidkhan_moh
WhatsApp	+39 331 106 0553
Address	Via Giuseppe Mazzini 4, 95032 Belpasso (CT), Italia

The flagship products

AI Tools

Product	Address	What it does
Claude Command Center	claude.cabscloud.com	AI operations & marketing command center.
Content Engine	contengine.pro	AI content generation at scale.
NexAI	nexai.cabscloud.com	AI content, chat, image, code & voice suite.
YouTube Pro	youtubepro.cabscloud.com	Grow your channel with AI tooling.
ReelForge	reelforge.cabscloud.com	UGC & faceless short-form video, generated.
AvatarTunes	tunes.cabscloud.com	AI avatar videos & voiceovers.
AgentLabs	agentlabs.cabscloud.com	AI voice-agent & call automation.
IntelliCall AI	intellicall.cabscloud.com	AI voice-calling agents & telephony.
AutoGineer AI	ai.autogineer.com	135+ AI tools: chat, content, code & voice.
Krikkit	krikkit.cabscloud.com	Describe a site, get a real one.

Marketing & SEO

Product	Address	What it does
LeadGen	lead.cabscloud.com	AI-powered lead generation & outreach.
66Audit	66audit.pro	AI SEO site audit & monitoring.
AtoZ SEO Tools	seo.codeengineers.net	Free all-in-one SEO tools suite.
WhatSaaS	whatsaas.cabscloud.com	WhatsApp business automation platform.
AutoGineer WhatsApp	whatsapp.autogineer.com	WhatsApp CRM, campaigns & team inbox.
The Tech ABC WhatsApp	whatsapp.thetechabc.com	WhatsApp Cloud API marketing SaaS.
whatsCRM	whatscrm.cabscloud.com	WhatsApp, Instagram & Telegram marketing.
Pitchbar	pitchbar.cabscloud.com	AI sales chat for high-intent pages.

Productivity

Product	Address	What it does
Organized Chaos	oc.cabscloud.com	Turn scattered notes into structured action.
Webby	webby.codeengineers.net	AI website builder & hosting.
BetterLearn	betterlearn.cabscloud.com	AI-powered learning & study tools.
AppForge	appforge.cabscloud.com	Turn any website into a native mobile app.
Autogineer Cloud	cloud.autogineer.com	Instant VPS hosting, deploy in seconds.

Media

Product	Address	What it does
StoresWiki Channels	channels.storeswiki.com	Every YouTube channel in the network, with live subscriber counts.

Working together

The work behind this dossier is also available on its own, for a buyer's team or for anyone who wants a system like it:

- **Autonomous content systems** built on the client's own cloud: channels, renderers, quality gates, publishing.
- **SaaS products** from an idea to a live multi-tenant platform: teams, credits, billing, APIs, documentation.
- **AI agents and automation**: Claude-based agents, n8n workflows, integrations with the platforms' APIs.
- **Cloud infrastructure**: GPU fleets on AWS and Google Cloud, cost control, backups and turn-key disaster recovery.

CONTACT
Majid Khan Mohmand
WhatsApp +39 331 106 0553 · linkedin.com/in/majidkhanmohmand · fiverr.com/majidkhan_moh
Via Giuseppe Mazzini 4, 95032 Belpasso (CT), Italia

REFERENCE

Appendices

The original programme's 65 channels, ChannelForge's channels, and a glossary of the terms this dossier uses.

APPENDIX A

Channel directory

All 65 channels of the original programme, by project. Handles of pre-launch channels are the planned names; their accounts are created at launch.

Channel	Proj.	Niche	YouTube	Facebook	Instagram	Status
Body Disclosures	1	the human body explained, inside out	@bodydisclosures	bodydisclosures	bodydisclosures	live
Space Simulated	1	space at true scale: journeys, orbits, stars and black holes	@spacesimulated	spacesimulated	spacesimulated1	live
History Simulated	1	history reconstructed: wonders, engineering and events	@historysimulated1	historysimulated	historysimulated	live
What if	1	what-if scenarios: disasters and physics played out	@whatilosophy	simulatedwhatifs	whatilosophy	live
Inside The Machine	2	how machines work, in 3D exploded views	@machineinsidious	machineinsidious	machineinsidious	live
Mind Tricked	2	optical illusions and brain teasers (photosensitive-safe)	@mindsparta	mindsparta	mind.sparta	live
Into The Deep	2	the deep sea and its life	@deependables	deependably	deependables	live
Oddly Simulated	2	oddly satisfying physics simulations	@oddlysimulations	oddlysimulated	oddlysimulations	live
Quiet Proof	3	visual proofs in mathematics	@quietproof	quietproof	quietproof	pre-launch
Colossal Scale	4	machines and structures at colossal scale	@colossalscales	colossalscales	colossal.scale	live
Model Cutaway	4	how AI and modern tech work, in cutaway	@ModelCutaway	modelcutaway	model.cutaway	live
Downwash	4	flight and aerodynamics	@downwashphysics	downwashphysics	downwashphysics	live
Fracture Atlas	4	why materials break: engineering failure	@fractureatlas	fractureatlas	fractureatlas	live
Unseen View	4	what machines sense that people cannot	@un.seenveew	unseenview	unseenview	live
Upright Robots	4	how robots move and balance	@UprightRobots	uprightrobots	uprightrobots	live
Signal Voyage	4	how signals travel	@signalvoyager	signalvoyage1	signalvoyage	live
In Phone Charges	4	the energy budget of everyday devices	@inphonecharges	inphonecharges	inphonecharges	live
The Quiet Cavern	5	caves and karst	@thequietcavern	thequietcavern	thequietcavern	pre-launch
Rivers of Ice	5	glaciers and ice	@RiversOfIce	riversofice	riversofice	pre-launch
Stepped Leader	5	lightning	@SteppedLeader	steppedleader	steppedleader	pre-launch
Millimetre Earth	5	the Earth at the millimetre scale	@MillimetreEarth	millimetreearth	millimetreearth	pre-launch
Slipface	5	dunes and moving ground	@slipface	slipface	slipface	pre-launch
Tephra Town	5	volcanoes and eruptions	@TephraTown	TephraTown	tephratown	pre-launch
Tremor Lines	5	earthquakes and plate motion	@TremorLines	tremorlines	tremorlines	pre-launch

Channel	Proj.	Niche	YouTube	Facebook	Instagram	Status
Into the Updraft	5	storms and the atmosphere	@intotheupdraft	intotheupdraft	intotheupdraft	pre-launch
Wild Blueprint	6	what animals build	@WildBlueprint	wildblueprint	wildblueprint	pre-launch
Built to Endure	6	survival in extreme places	@BuiltToEndure	BuiltToEndure	builttoendure	pre-launch
Deep Time Street	6	deep time and fossils	@DeepTimeStreet	deeptimestreet	deeptimestreet	pre-launch
Homeward Atlas	6	animal navigation and migration	@HomewardAtlas	HomewardAtlas	homewardatlas	pre-launch
Pursuit Lines	6	predators and chases	@PursuitLines	pursuitlines	pursuitlines	pre-launch
Wild Sensorium	6	animal senses	@WildSensorium	wildsensorium	wildsensorium	pre-launch
Smaller Than Rain	6	life smaller than a raindrop	@SmallerThanRain	smallerthanrain	smallerthanrain	pre-launch
Swarm Logic	6	collective behaviour in animals	@SwarmLogic	SwarmLogic	swarmlogic	pre-launch
Beforehand	7	how it was done before the app	@Beforehand	Beforehand	beforehand	pre-launch
Why That Size	7	why things are the size they are	@whythatsize	whythatsize	whythatsize	pre-launch
Cofferdam	7	megaprojects, built step by step	@cofferdam	cofferdam	cofferdam	pre-launch
Below The Block	7	the hidden city under the street	@belowtheblock	belowtheblock	belowtheblock	pre-launch
Domino Trail	7	chain reactions and ripple effects	@DominoTrail	DominoTrail	dominotrail	pre-launch
Odd Ledger	7	price anomalies and odd economics	@OddLedger	oddledger	oddledger	pre-launch
The Load Path	7	how structures carry their loads	@theloadpath	theloadpath	theloadpath	pre-launch
Throughput Lab	7	how crowds, queues and traffic flow	@throughputlab	throughputlab	throughputlab	pre-launch
Brake Run	8	the engineering of thrill rides	@brakerun	brakerun	brakerun	pre-launch
Engineless	8	wind-powered craft	@engineless	engineless	engineless	pre-launch
Halfspace	8	tactics, drawn on the board	@halfspace	halfspace	halfspace	pre-launch
Lean Angle	8	the physics of two wheels	@leanangle	leanangle	leanangle	pre-launch
Line Call	8	photo finishes and line calls	@linecall	linecall	linecall	pre-launch
Midflight	8	the physics of sport	@midflight	midflight	midflight	pre-launch
Railhead	8	rail technology	@Railhead	Railhead	railhead	pre-launch
Tenths	8	motorsport technology	@tenths	tenths	tenths	pre-launch
Endgrain	9	woodcraft and timber	@endgrain	endgrain	endgrain	pre-launch
Heatwork	9	crafts of fire: glass, metal, ceramics	@heatwork	heatwork	heatwork	pre-launch
Within a Micron	9	precision making	@withinamicron	withinamicron	withinamicron	pre-launch
Off The Line	9	how production lines make things	@offtheline	offtheline	offtheline	pre-launch
Under the Patina	9	the science of restoration	@underthepatina	underthepatina	underthepatina	pre-launch
Specimen Bench	9	wonder materials	@specimenbench	specimenbench	specimenbench	pre-launch

Channel	Proj.	Niche	YouTube	Facebook	Instagram	Status
Unbuilt	9	the science of demolition	@unbuilt	unbuilt	unbuilt	pre-launch
Unfurled	9	engineering that folds	@unfurled	unfurled	unfurled	pre-launch
Envenom	10	the science of venom	@envenom	envenom	envenom	pre-launch
Fledge	10	how animals raise their young	@fledge	fledge	fledge	pre-launch
Host Hijack	10	parasites that steer their hosts	@hosthijack	hosthijack	hosthijack	pre-launch
Linchpin	10	food webs and keystone species	@linchpinwild	linchpinwild	linchpinwild	pre-launch
Mutual Terms	10	partnerships in nature	@mutualterms	mutualterms	mutualterms	pre-launch
Mycelium Minute	10	the fungal world	@myceliumminute	myceliumminute	myceliumminute	pre-launch
Outlived	10	the longest-lived living things	@outlived	outlived	outlived	pre-launch
Now You Don't	10	camouflage	@nowyoudont	nowyoudont	nowyoudont	pre-launch

APPENDIX B

ChannelForge's channels

The channels on ChannelForge today: the owner's own Enterprise workspace (Khan Studio), which runs 43 channels in 6 projects, most at two episodes a day each, and the quality-assurance workspace used to test every release. Most of the newest have no platform accounts yet: they make episodes into their stock and post once connected.

Channel	Project	Niche	State	Per day	Episodes
Everyday Workings	QA workspace	How machines work	ready	0	10
Steam and Grounds	QA workspace	How Moka Express 3-cup works	ready	0	3
Load Path	QA workspace	The Forces Inside Buildings	ready	0	3
Inside The Machine	QA workspace	How machines work, in exploded view	ready	0	4
Inside The Machine · Japanese	QA workspace	How machines work, in exploded view	ready	0	1
Airside	Built for people	The Airport Behind the Gates	ready	2	3
Coolth	Built for people	Homes That Beat the Weather	brand	2	3
Freeboard	Built for people	Inside a Working Ship	ready	2	2
Hermetic	Built for people	Living Where People Shouldn't	ready	2	2
Scrim	Built for people	How Live Shows Are Engineered	ready	2	3
Superseded	Built for people	How Old Gadgets Really Worked	brand	2	2
Unscathed	Built for people	How Deadly Forces Become Survivable	ready	2	2
Windrow	Built for people	How Modern Farms Work	ready	2	3
Antinode	Custom niches	The Physics of Sound	ready	2	45
Fond	Custom niches	The Chemistry of Cooking	ready	2	47
Footfall	Custom niches	How Animals Move	ready	2	45
Hairline	Custom niches	Engineering Failures, Explained	ready	2	46
Load Bearing	Custom niches	How Structures Stand Up	ready	2	45
Slipface	Custom niches	How the Ground Moves	ready	2	45
Turgor	Custom niches	Plants Are Doing Things	ready	2	46
Virga	Custom niches	What the Sky Is Doing	ready	2	43
Allotrope	Everyday wonders of science	One Element at a Time	ready	2	3
Chronoscope	Everyday wonders of science	Fast Events, Slowed Down	ready	2	3
Druse	Everyday wonders of science	How Crystals Grow	ready	2	3
Eigenstate	Everyday wonders of science	The Quantum World, Visible	ready	2	3
Halation	Everyday wonders of science	How Cameras Catch Light	ready	2	3
Lightfast	Everyday wonders of science	Why Things Are That Colour	ready	2	3
Prior Art	Everyday wonders of science	Inventions Stolen From Nature	ready	2	3
Shoaling	Everyday wonders of science	What Waves Do to Coasts	ready	2	3

Channel	Project	Niche	State	Per day	Episodes
Boustrophedon	Minds, games and big ideas	Cracking Lost Writing and Codes	ready	2	3
Gedanken Lab	Minds, games and big ideas	Famous Thought Experiments, Made Visible	ready	2	3
Kenning	Minds, games and big ideas	Great Myths of Ancient Worlds	ready	2	3
Least Count	Minds, games and big ideas	The Experiments That Changed Science	ready	2	3
Misjudged	Minds, games and big ideas	Why Smart People Get Fooled	ready	2	3
Morphospace	Minds, games and big ideas	Designing Life for Other Worlds	ready	2	3
Speciate	Minds, games and big ideas	Animals Ten Million Years From Now	ready	2	3
Zwischenzug	Minds, games and big ideas	Chess Made Visible in 3D	ready	2	3
Afterimage	Proven formats	Illusions and brain teasers	ready	2	46
Erstwhile	Proven formats	Be there in history	ready	2	44
Fathom Down	Proven formats	Journeys into the deep	ready	2	45
Lightminute	Proven formats	Space, to scale	ready	2	48
Otherwise Earth	Proven formats	What if…?	ready	2	48
Settle	Proven formats	Oddly satisfying physics	ready	2	45
Unbolted	Proven formats	How machines work	live	3	62
Vitreous	Proven formats	What happens to your body	ready	2	44
Space Simulated	The original's channels on TikTok	space	ready	0	0
History Simulated	The original's channels on TikTok	history	ready	0	0
What if	The original's channels on TikTok	what if	ready	0	0

APPENDIX C

Glossary

Term	Meaning
Agent	an AI model working in a sandbox with tools: it writes, runs and checks code
Batch	a set of episodes made together on one GPU box for one project
Blender episode	an episode whose picture is a 3D scene built in Blender from free CC0 assets
Caption sync	how far each caption is from its spoken word, measured on the final soundtrack
Catch-up	an extra post or generation that makes up for a missed slot
CC0	a public-domain dedication: an asset anyone may use for anything, with no attribution owed
Critic	the model that scores an episode out of 10 before it may be posted
Edit (90-second)	the ≤ 89-second cut posted to Facebook and Instagram
Failure ledger	the append-only record of every failed publishing attempt and how it ended
Flag	a switch that turns a feature on for every workspace, for some, or for none
God page	a page grown to show everything at once; split into tabs since November
Golden image	a machine image every GPU box is made from, so all are identical
In step	every platform of a channel has posted the same episodes
Kit	a channel's renderer: its scene library, style rules and references
L4	NVIDIA's 24 GB data-centre GPU used for rendering and agents
Lever	the button beside a problem that fixes or retries it, where the problem is shown
Owed	a post a slot could not make, to be made up from surplus
Pool	the GPU boxes one project or lane may use, tried in order
Posting guard	a check of the platform's own posts before every publish
Pre-overhaul	an episode made before the video format's overhaul of 22 October; held at launch
Reconciler	the program that turns a pool's intent into running or stopped boxes
Release floor	the bar a held film must reach to be posted: 8/10 overall, 7 or more per act
Sentinel	the five-minute observer that acts on whatever is wrong
Slot	a scheduled posting time of a channel on a platform
Studio	the portal's button-driven builder for projects, channels, kits and brands
Word timing	the time of each spoken word, heard back by a speech recogniser after the voice is made
Yield	pre-launch work giving its GPUs to a live project's waiting batch

THANK YOU

For questions, a demonstration of either system, or an offer: Majid Khan Mohmand, WhatsApp +39 331 106 0553, linkedin.com/in/majidkhanmohmand, fiverr.com/majidkhan_moh. This dossier is also a website, with every chart live: dossier.storeswiki.com.