Finish one film first.
Then learn to create.

TimeGleam · AI Video Creation & Distribution

Start with: The Last Lamp exercise (workflow / free canvas / MCP Agent)
Then unpack: two short-story walkthroughs (Don't Knock, Paper Boat)
Then learn: seven stages / eight lessons / six craft chapters
Desktop app and remote MCP / canvas manual / 85 tools
Full cases: five films

Edition 2026-11-14
Same source as the public web edition, with examples, screenshots and chapter bookmarks.
Pick one case first; you do not have to read the manual end to end.

Contents

01 Start with the Result
Real-world reference and the shared exercise 6

Reading guide and exercise production . . . 6
Case 1: workflow "one-click works" - The Last Lamp, finding the keychain, one 8-second shot . . . 9
Case 2: free canvas "one-click film" - The Last Lamp, six shots, about 30 seconds . . . 11
Case 3: MCP Agent - continue S04 in The Last Lamp . . . 15
Optional: make one shot manually from one image . . . 17

02 Story Drills
Two ten-second-plus films, walked from story to finished film 18

Drill 1 · Don't Knock: one action carries 16 seconds . . . 18
Drill 2 · Little Boat Across the Puddle: no characters, but still a protagonist . . . 23
From story to finished film: seven stages, and what AI and prompts handle . . . 28

03 Eight Starter Lessons
the shortest manual path 31

Get to know the creation space . . . 32
Start with an idea . . . 34
Prepare assets and subjects . . . 37
Make the first frame . . . 40
Bring the FRAME to life . . . 42
Finish and fix the other shots . . . 44
Editing and sound . . . 47
Export and continue creating . . . 49

04 Finished-Film Cases
Scripts, shot lists and layered production for five full short films 51

Case 4: A complete production · from script to The City of Spare Charge . . . 51
Case 5: A complete nature-fantasy film · from script to The Spirit of Spring Mountain . . . 71
Case 6: A complete snow-epic production - from script to The Bell on the Snow Ridge . . . 94
Case 7: Complete dual-line narrative film · From script to "Cloud Ridge Letter" . . . 116
Case 8: Full-flow rebuild · from script breakdown to "Snowridge Bells 2 · One Bell Returns" . . . 151

05 Advanced creation
story, shots, continuity, colour, edit and review 169

Advanced creation: Make better videos . . . 170

Advanced 1: Story and emotion . . . 171
Advanced 2: Storyboards and visual language . . . 173
Advanced 3: Continuity and consistency . . . 175
Advanced 4: Light, color and art direction . . . 177
Advanced 5: Editing rhythm and sound . . . 179
Advanced 6: Review and targeted fixes . . . 181
Complete film plan: three scenes, twelve shots . . . 184

06 Desktop app, MCP, and automation
plug a local Agent into the same creative space 189

TimeGleam Desktop: put an Agent workbench in the same creative space . . . 189
MCP hands-on: continue in the same work . . . 192
Connect an AI tool . . . 192
MCP: First understand IDs, revisions and receipts . . . 193
Walkthrough 1: Start in the web app, edit with MCP, continue in the web app . . . 194
Hands-on 2: Check inputs and costs, and edit only the target shot . . . 195
Continue the same film across clients . . . 197
Walkthrough 4: Adopt free content into the official workflow . . . 199
Walkthrough 5: Lost responses, 409, and interrupted jobs . . . 200
MCP integration, three guides, and Agent prompts . . . 201

07 Freeform Canvas Manual
The complete manual, organised by feature 202

Complete feature manual for the free canvas . . . 203
Canvas manual: documents and nine object types . . . 203
Finding works: My works and Works in this workspace . . . 205
Asset reuse: My assets and cross-workspace management . . . 208
Canvas guide: navigation, selection, and shortcuts . . . 210
Canvas guide: sections, groups, content, and templates . . . 211
Canvas guide: models, specs, and real inputs . . . 212
Canvas guide: text candidates and shot drafts . . . 214
Canvas guide: assets, subjects, and formal adoption . . . 216
Canvas handbook: all image tools . . . 217
Canvas manual: voiceover, audio, and video sampling . . . 218
Canvas manual: frame-by-frame review . . . 219

Canvas guide: Panorama and 3D director stage . . . 221
Canvas guide: candidates, tasks, and three histories . . . 223
Canvas guide: Assistant and batch production . . . 224
Canvas guide: One-click film pricing and independent editing . . . 226
Canvas guide: editing, subtitles, checks, and delivery . . . 230
Canvas guide: saving, permissions, mobile, and cleanup . . . 231

08 More Tools and Troubleshooting
Workflows, adaptation, extension tools and common questions 234

Workflow guide: final scripts and batch production . . . 234
Start production from STORY . . . 237
More creation tools . . . 237
If you hit a problem, continue here . . . 238

09 Tool Index
The full remote tool contract 240

MCP: complete index of 85 tools . . . 241

01 Start with the Result
Real-world reference and the shared exercise

Reading guide and exercise production

TimeGleam · AI video creation and distribution

Tutorial checked: 2026-11-14 · starter, drills, advanced craft, and the full manual

GATHER THE LIGHT. GET IT SEEN. From one brief to a playable short film.

The tutorial has nine parts. The shared exercise is The Last Lamp: a rainy night, an old grocery store, a worn brass keychain engraved with 回家. It starts as one 8-second shot in case 1, grows into a six-shot story of about 30 seconds in case 2, and is then handed to the Agent for a targeted revision in case 3. The real-world reference is Shift Change in the Happy Story space of the Happy Kai account, used to check what a five-shot chain looks like when it really ships. In between, two short stories, Don't Knock and Little Boat Across the Puddle, take story, script, character, scene, prop, storyboard and finished film apart step by step, and show which part AI and prompts control at each stage. To start making, go to Eight lessons; to put a local Agent to work, read TimeGleam Desktop and MCP in practice.

Choose one method first and produce your first result

What do you want to do first?	Start here
Start with a real finished film, to see that a chain really ends in a playable file	Shift Change · 19.25 seconds · five shots
Follow one shared exercise from an 8-second shot to a six-shot film of about 30 seconds	Case 1 · finding the keychain → Case 2 · the complete six shots
See how much of story, script, character, scene, prop, storyboard and finished film goes to AI	Don't Knock · 16 seconds · Little Boat Across the Puddle · 15 seconds · Seven stages
Run one short shot automatically in workflow and see which work it lands in	Case 1 · 8-second single shot

What do you want to do first?	Start here
Let the canvas plan six shots, then check, adopt and export them one by one	Case 2 · 30 seconds, six shots
Let an MCP-connected Agent continue in the same The Last Lamp	Case 3 · continue S04
You have the desktop app and want local files, scripts, and an Agent in the same workspace	TimeGleam Desktop
Watch a 75-second industrial sci-fi film and follow the full path: script, shot list, previz, film, edit, export	Case 4 · The City of Spare Charge
Watch an 82-second nature-fantasy film and practise layered production with several characters, an animal, a stone spirit and four seasons	Case 5 · The Spirit of Spring Mountain
Watch an 87-second snow-epic film and practise few characters, hard light, animal fur and a prop turntable	Case 6 · The Bell on the Snow Ridge
Watch an 81-second Cloud Sea Postal Route film and practise dual-line narrative with one letter running through the whole film	Case 7 · Cloud Ridge Letter
Walk the whole path once: script breakdown, 30 asset cards, 12-shot previz, layered rendering, AI retouch and platform assembly	Case 8 · Full-flow rebuild

Learn and practice over several sessions. Asset generation time varies with the model and queue and is not guaranteed to take a fixed time. The tutorial is publicly readable; actual creation requires sign-in, the relevant workspace permissions, and available models.

Shared exercise: "The Last Lamp"

On a rainy night a young man walks into a grocery store that is about to close and finds a worn brass keychain engraved with 回家 on the wooden counter. He hesitates, reaches out, stops, then takes it, turns off the last lamp over the counter and walks into the rain.

Film target	About 30 seconds
Aspect ratio	16:9
Scale	1 person, 1 location, 1 prop
Shot	6 shots

This is an exercise written for the tutorial. It is not a finished film already published in the Happy Story space, and it does not reuse another production's media. It keeps one character, Lin Mo, one old grocery store and one keychain so that every stage stays visible: the story has a turn, the script has filmable actions, the character and location have clear identifying details, the prop carries the 回家 information, and the storyboard can be checked shot by shot.

The real-world reference is Shift Change in the same space: 19.25 seconds, five shots, no dialogue, already exported and playable. It proves a five-shot chain can be completed, but the production is mixed: the first three shots use previz footage directly in the cut, the last two come from still images turned into video. The tutorial does not retell its story; it only uses it as a comparison for what a real delivery looks like.

The two short-story walkthroughs do not repeat these two productions. They use Don't Knock and Little Boat Across the Puddle: 16 and 15 seconds, three to four shots, built specifically to rehearse the chain from story to script, character, scene, prop, storyboard and finished film, and to show what AI and prompts handle at each stage. The three opening cases then return to The Last Lamp: case 1 makes one 8-second shot in the workflow, case 2 covers six shots in the free canvas, case 3 has the Agent change only S04.

Examples 4 to 8 are five longer self-made films: a 75-second industrial sci-fi piece for layering and assembly, an 82-second nature fantasy for multiple characters and four seasons, an 87-second snow epic for hard light and animal fur, an 81-second cloud-sea postal route for dual-line narrative and vehicle continuity, and a final one that rebuilds the whole path from script breakdown to platform assembly. The eight lessons below go deeper, step by step.

> Scope note: 30 seconds is the target for the exercise film; a new workflow production defaults to 60 seconds, so change it to 30 for the exercise. Pick a per-shot length the model supports and adjust it on the timeline. The real specs and the mixed production of Shift Change are a reference only and must not be read as proof that six shots were generated automatically in one run.

Case 1: workflow "one-click works" - The Last Lamp, finding the keychain, one 8-second shot

What you see first: One 8-second shot: on a rainy night in an old grocery store, Lin Mo finds a worn brass keychain on the counter under the warm lamp, reaches out slowly, and stops his fingertips just before touching it. The system runs script, production prep, shot video and the first cut on its own, and saves the result as a new production.

Studio home → Workflow creation → Quick work

Why start from this shot? It holds a single continuous action yet tests four things at once: whether the character stays consistent, whether the keychain reads at a glance, whether the hand meets the counter naturally, and whether the warm store lamp and the deep-blue rainy night sit in the same space. Get any one wrong and the following six shots drift with it. It also shows the edge of one-click workflow best: the system strings together script, preparation, generation and the first cut, but which shot to make, which take to accept and whether the engraving is credible are still human calls.

This shot is not a fragment of Shift Change or any finished piece. It is the starting point of the shared exercise The Last Lamp. Make the 8 seconds first, confirm that the character, the prop and the light hold up, then expand it into six shots in case 2.

Copy this brief

The Last Lamp, first-shot test. Rainy night in an old grocery store; Lin Mo, 25, in a dark grey old jacket, finds a worn brass keychain on the wooden counter, engraved with 回家 on the back. He reaches out slowly with his right hand and stops his fingertips just before touching it, without picking it up or turning away. One continuous shot, camera pushing in slowly; warm yellow store light, deep blue rainy night outside, realistic film look. No dialogue, no subtitles, no music, natural ambience only. The character's clothing, the keychain and the store interior stay consistent throughout, and no second customer appears in frame.

Follow along

1. **Open Quick work and paste the brief:** Paste the text above into Content, idea or script. Choose 8 seconds and 16:9; pick a resolution the current model allows. The 5/8/12-second options are what the workflow offers, not the 30-second target of case 2 and not a fixed length for the eight lessons. A new workflow production defaults to 60 seconds; this case changes it to 8 manually.

2. **Review settings, then confirm and start:** In More settings · Models and checkpoints, confirm the text, image, and video models and keep "cut and export a finished film automatically when the shot completes". For a first fully automatic run, leave the manual checkpoint unselected; to see the script first, tick the matching pause. When the checks pass, click "Confirm and start production"; this creates a new work and submits billable model tasks.

3. **Wait on the same production record:** Watch the status bar and Production details; you can leave the page and reopen the run from One-click works on the home page. If a manual checkpoint appears, check the character, the keychain and the action before continuing. When something needs attention, read the reason on the original task instead of creating a second production from the home page.

4. **Play the first cut:** When it finishes, click View first cut, or inspect shots through Video preview. Save the file and play it; turning auto-edit off only stops at video preview and is not an exported film.

Completion criteria: A playable first-cut file in which you can see Lin Mo's dark grey jacket, the brass colour of the keychain and the 回家 engraving, and the moment the fingertip stops; a generated script is not the finish line. Sound and picture still need your own listening and checking, and no model is guaranteed to satisfy the brief every time.

Text, image, video, and other model calls are billed by actual tasks; Workflow "Quick work" is not the canvas "Generate plan and price" flow and will not first give a single whole-film confirmation sheet. Pausing only stops further progress; submitted tasks may still complete. If the response is unclear, first use "Recover last work."

To fix this production, read Workflow guide; to expand this shot into a full story, continue with Case 2.

Case 2: free canvas "one-click film" - The Last Lamp, six shots, about 30 seconds

What you see first: Take one moment from case 1 and open it into a full story: rainy street, pushing the door open, walking along the counter, finding the keychain, fingertips over the engraving, turning off the light and walking into the night. The canvas keeps the plan, the production record and the film node, and the film has its own edit and export.

Studio home → Canvas creation → One-click canvas film → Generate plan and price → Start film

Case 1 proves one action can be generated; case 2 proves six actions can form a single choice. The difficulty is no longer a single frame but planning: how the six shots divide the work, which shot carries the information, in which shots the keychain must be legible, and whether the character and the store keep a consistent direction. Asking AI for a plan first is exactly how these questions surface in text before money is spent; a person must check each shot before confirming the quote, not after all six clips are generated.

Six-shot plan (about 30 seconds)

The Last Lamp teaching shot list (6 shots / 30 seconds, ready to paste into shot drafts)

Shot	Duration	Shot size / camera	Image and action	What the shot gives the audience
S01	4 s	Wide, locked off	Rainy street; only the counter lamp of the old grocery store is still on; Lin Mo walks in from frame right under a black umbrella	Time, weather, place, and "the last lamp"
S02	4 s	Medium, eye level	At the door he closes the umbrella, shakes off the water and pushes the door open; the bell chimes softly	The character enters the space; sound arrives before any dialogue
S03	5 s	Wide, lateral track	Shelves, pendant lamp and an old radio enter frame in turn as he walks along the counter	Store layout and the character's position, which set the direction for later shots
S04	5 s	High angle, slow push-in	A worn brass keychain on the wooden counter, engraved with 回家 on the back	Reveals the prop; the engraving is the turning information

Shot	Duration	Shot size / camera	Image and action	What the shot gives the audience
S05	6 s	Close shot into hand close-up	He stops, reaches out with his right hand, fingertips hovering over the engraving, not picking it up yet	Hesitation and choice; the hand is the centre of the performance
S06	6 s	Medium to wide	He closes the keychain in his palm, turns off the counter lamp, pushes the door open and walks into the rain	Taking it away and saying goodbye, ending on the aftertaste of 回家

The six shots total 30 seconds, each with one main action. S01 and S03 establish space, S02 brings the character in, S04 delivers the information, S05 the choice, S06 the result. Leave a quarter second of margin at each cut; if the model can only make a shorter clip, fill the gap on the timeline instead of packing three actions into one prompt.

Copy this brief

Make The Last Lamp, about 30 seconds, 16:9, realistic film look, no dialogue, no subtitles, no on-screen text.
Rainy night; Lin Mo, 25, in a dark grey old jacket, walks into a grocery store that is about to close and finds a worn brass keychain on the wooden counter, engraved with 回家 on the back. After hesitating he takes it, turns off the last lamp over the counter and walks into the rain.
Six shots, one main action each:
S01 wide, rainy street, only the counter lamp of the old grocery store still on, Lin Mo walks in from frame right under a black umbrella;
S02 medium, at the door he closes the umbrella, shakes off the water and pushes the door open, the bell chimes softly;
S03 wide lateral track, shelves, pendant lamp and an old radio enter frame in turn as he walks along the counter;
S04 high angle slow push-in, the worn brass keychain on the wooden counter, engraved 回家 on the back;
S05 close shot into a hand close-up, he stops, reaches out with his right hand, fingertips hovering over the engraving;
S06 medium to wide, he closes the keychain in his palm, turns off the counter lamp, pushes the door open and walks into the rain.
Shared constraints: the same character, the same old grocery store, the same worn brass keychain engraved with 回家; warm yellow store light against the deep blue rainy night; clothing, prop appearance and store layout consistent throughout; no second customer in frame.

Follow along

1. **Write the brief in the canvas creation bar:** Paste the text above, use "Add references" as needed, then click "One-click canvas film." "New blank canvas" and the manual entries under "More ways to create" only create a draft and do not run automatic production.

2. **Choose model specs, generate the plan and price it:** Check the authorized video model, the video resolution, 16:9 and the 30-second target. More settings lets you choose the plan text, image and voice models, the per-shot length, the reference mode and the export resolution. Click Generate the plan and price it to run the billable planning step first; images, video and audio are not produced yet.

3. **Check all six shots, then confirm:** When the plan is ready, check shot by shot: does S01 show only one lamp left, can the engraving in S04 be read, does the hand stop in the right place in S05, does S06 actually turn the light off. Check the shot count, the items to generate and the items to reuse; the settled planning charge is listed separately and is not counted again in the later hold. Confirm that the hold shown for this run is acceptable, then click Start the film. The home page does not take a manual point cap and does not authorise spending the whole balance; sample points are not a quote.

4. **Review and export from the original record:** On the canvas, check progress, cost, and pause reasons for up to 8 shots, then open the film's own edit to check order, rhythm, and sound. Download the complete file and play it from the start; the file defines the actual duration. If it stops, return to "Continue the saved canvas plan · do not replan" on the home page or the original canvas; do not create another work.

Completion criteria: A playable six-shot result inside the same film node, a complete downloadable file, and a clear read of the choice from discovery to hesitation, taking it and turning off the light; a plan, a single shot or an export queued still does not count as done.

Generating only the plan also calls a model and costs money; it is not a free preview. The finished film is an independent edit and is not written into the formal episode storyboard or the main timeline. This case did not run a real paid generation in the Happy Story space; sample points and specs are not a quote.

Real-world reference: Shift Change



If it will not play, download the 480p film (854×480, 19.250 seconds, 24 fps, 1.0 MB).

A real finished film in the Happy Story space of the Happy Kai account: 19.25 seconds / 5 shots / 854x480 / 24 fps / ambient sound. Watch three contact points as it plays: the wall clock jumping to 04:00, the door sliding open by itself, and a hand pushing the cat-food tin across.

This film is a real export, but its production is mixed: s0-s2 use previz footage directly in the cut, s3-s4 come from still images turned into video. It proves that a five-shot chain can run through and export, not that all five shots came out of one automatic run. It sits here so that, while you plan six shots, you have a delivery standard to check against: the file plays, the actions connect, the sound lands, and it is more than a few pretty frames.

Quote, locking and continuation details are in One-click film handbook. To let the Agent operate, continue with Case 3; to learn to fix shots yourself, start from Eight-lesson primer.

Case 3: MCP Agent - continue S04 in The Last Lamp

What you see first: without creating anything new. The Agent first reads the The Last Lamp you already saved: six shots, current takes, timeline references and export records; then it changes only the S04 you name, re-quotes, waits on the original job and exports to a new file. The web page can then be reopened to check the result, rather than a script handed over in chat.

An Agent client that supports MCP → connect to TimeGleam MCP → send the case prompt

First complete browser authorization with Connection and authorization and MCP setup and confirm the client can read TimeGleam tools. Without an Agent, read the first two web examples; there is no need to create or copy a key by hand.

Why S04? It carries the most information in the whole film: the keychain has to be in frame, the engraving has to be readable, and the camera still has to push in slowly from outside the counter. It is a good test of whether the Agent really read the current take and the timeline instead of regenerating from the start. It also means you never paste a key into the chat and the Agent never guesses at cost.

Copy for the Agent

Use the connected TimeGleam MCP to continue The Last Lamp that I saved as the shared exercise in the Happy Story space. Do not create a new production.
Current state: about 30 seconds, 16:9, 6 shots (S01 rainy street / S02 pushing the door open / S03 walking along the counter / S04 keychain on the counter / S05 fingertips over the engraving / S06 turning off the light and walking into the rain), no dialogue, no subtitles.
Goal: change only S04. Keep Lin Mo's dark grey old jacket, the warm yellow store light, the wooden counter and the worn brass keychain; push the camera in slowly from outside the counter and end on the 回家 engraving on the back, steady and readable, with no picking up or leaving, still 5 seconds.
Read only first: locate the real production, episode and free canvas, read back the node ID of S04, the current take and the timeline reference, and tell me exactly what you will change and why the other five shots stay untouched.
Then give me the shot plan, the models used and whatever cost information is available, and wait for me to confirm scope and budget before calling any paid generation. If cost cannot be established, stop and say so instead of guessing.

Once I confirm, complete the S04 generation, the candidate review and explicit adoption, the replacement of the affected timeline instance and one new export within that scope, then give me a limited preview, the actual specs and a direct download link. Do not overwrite the existing production, the existing exports or the other five shots, and do not call web features that are not open yet.
Re-read the revision before every write and preview as the tools require; keep request_id, submission_id, impact_token and job_id, and after a network failure check the receipt and the original task instead of resubmitting a paid job. Pause on conflicts, anything beyond the authorised scope, or anything that needs human judgement.
Finally give me the production and episode IDs, the shots actually changed, the export job and asset IDs, the preview range and a direct download link, plus anything unfinished and how to continue. If there is no successful export, say plainly that it is unfinished.

Follow along

1. **Send the brief, review the plan first:** Send the text above to the Agent. Check the space, models and shots it reports, and confirm it identified S04 rather than starting again from S01.

2. **Confirm the production scope for this round in one go:** Give the cost range you accept, then authorize execution. film_quote stores the server quote; film_confirm must carry delegated, the original quote_id, quote_fingerprint, and the authorized amount after your explicit consent. An unavailable cost must not be treated as zero. Stop and come back to you if the price is out of range, unknown, or conflicting.

3. **Let it execute and keep a continuation record:** The Agent uses film_* to continue the same film; there is still no "Quick work" MCP interface in workflow. It waits for the original job, and candidate checks, adoption, timeline updates, and export each finish separately. If the client disconnects, a submitted background job is not cancelled, and Agent orchestration is not guaranteed to continue on its own.

4. **Get the files; refine in the web app if needed:** The Agent uses media_delivery to return the cover, actual specs, and a five-minute direct download link, checked against the work, episode, and film_id. A first frame is not the whole film; download it when the host cannot play it. If export is unfinished, check the original job first; after an interruption, give the Agent the record without credentials and follow the resume guidance instead of redoing everything.

Completion criteria: The new export downloads, the push-in and engraving in S04 really changed, the other five shots and the old exports are still there; "export started" or a confident summary is not completion. Results and timing depend on the model, the media and the queue, and this case did not run a real paid generation.

After the three opening examples, use Eight-lesson primer to learn how to adjust characters, shots, and sound yourself; or start with Manual single-shot practice below. External tools and interruption handling are covered in MCP in practice.

Optional: make one shot manually from one image

Prepare an image you have the right to use; if you do not have one, start with Lesson 4.

1. **Go to draft:** On the home page, add a reference image, then click Video and check the input reference; you can also upload from within the video node.

2. **Describe a simple action:** Choose the model and compliant specs, then write one action. Creating a draft doesn't generate; only clicking Generate submits.

3. **Preview and adopt:** Wait for the original task, check the result in Assets and Versions, and adopt it explicitly once it fits. Previewing does not change what is currently adopted.

4. **Download asset:** More asset actions → Download asset. Downloading a single shot is not the same as exporting the full timeline.

> Lin Mo looks at the old brass keychain on the counter, gaze slowly dropping, body leaning slightly forward. The camera pushes in slowly, the lighting holds steady, rain runs down the window. Keep the character, clothing, and space consistent; show only this one action.

> If a task fails or is waiting to recover, read the original task's reason first; don't keep creating new paid tasks.

Try again in your studio

02 Story Drills
Two ten-second-plus films, walked from story to finished film

Drill 1 • Don't Knock: one action carries 16 seconds

This is the shortest complete path in the tutorial: 16 seconds, 4 shots, one character, one prop, one location. Its value is not visual complexity but that every step is small enough to finish in one sitting, and still exposes real problems: hand action, light through the door gap, the contact point between cat paw and bag.



If it will not play, download the 480p film (854×480, 16.000 s, 24 fps, 1.5 MB).

Finished film: 16 seconds / 4 shots / 854×480 / 24 fps / ambience. While watching, note three joins: setting the bag down, pulling the hand back, the cat paw appearing.

Step 1: story - find the action that does not happen

One-line story: on a rainy night a delivery rider sets the food at the door, raises a hand to knock, pulls it back and leaves; a cat paw reaches out of the door gap and drags the bag inside.

This story has one turn, and the turn is exactly **not knocking**. Written as "a rider delivers food, sets it down and leaves" it is only a procedure; written as "he wants to knock but does not", the audience asks why. No dialogue supplies the answer; the cat paw in the door gap does: someone is home (or a cat is), and they should not be disturbed.

Protagonist	Delivery rider, mid-20s, raincoat soaked through
Want	Deliver the food and confirm someone got it
Obstacle	No answer from inside, only a little warm light and movement
Choice	Do not knock; set it down and go
Result	The cat drags the food inside; the kindness lands

The usual mistake when writing the story is to write emotion as action: "he sets the food down gently". Gentleness is not an action; the audience cannot see it. What can be shot is: crouch, straighten the bag, raise the hand and stop short of the door, curl the fingers back, stand up. Write those actions in the script stage and every later shot has something to stand on.

Step 2: script - three lines still have to be shootable

> Short Story · Don't Knock - about 16 seconds, 16:9, realistic film look, rainy night.
> 1. Night, corridor of an old residential block. A delivery rider in a dark raincoat carries a food bag to the door, walking softly.
> 2. He crouches at the door, straightens the bag, raises a hand to knock, and his knuckles stop a few centimetres short of the door.
> 3. Warm yellow light shows in the door gap; there is a faint movement inside. He waits two seconds, withdraws his hand, stands and leaves.
> 4. Only the bag is left at the door. A ginger cat's paw reaches out of the gap, hooks the bag opening, and drags it inside a little at a time.
> No dialogue, no subtitles, no score; keep only rain, footsteps and the rustle of the bag.

Note three things: first, state the **light** (cold corridor light, warm light in the door gap), which sets the palette of every later shot; second, state the **Contact points** (hand stopping at the door, cat paw hooking the bag opening), which is what AI most often draws wrong; third, state the **what to leave out** (no dialogue, no subtitles, no score) to stop the model improvising.

Step 3: character, location, prop - three cards that lock the look

Three asset cards, each owning one thing

Asset	What to write	Why
Character · delivery rider	Male, mid-20s, short black hair, dark grey raincoat with a reflective strip across the chest, black rain boots, bag in the right hand	The raincoat and reflective strip are the only readable silhouette; all four shots must match
Location · corridor	Old residential block, beige walls with water stains, chipped paint on the steel door, a crooked door plate, one cold white lamp at the end of the corridor	It fixes the spatial relations and the light direction, so every shot does not invent a new corridor
Prop · food bag	Kraft paper bag, white handles, one small grease stain, the opening slightly apart	The opening must be hookable by a cat paw; the grease stain is the identifying detail

The three cards do not need to be long. The test is: can someone else draw the same character, the same door and the same bag from the card alone? "A handsome young man" or "a shabby but atmospheric corridor" cannot be checked, so they count as nothing written.

In TimeGleam these three go into subject packs: front and profile for the character, wide shot and door close-up for the location, front and bag-opening detail for the prop. The main image fixes the look and the others add detail; auxiliary images do not become the official version by themselves, and adopting one is an explicit action on the node.

Step 4: shot list - each of four shots carries one piece of information

Shot list (16 seconds / 4 shots)

Shot	Duration	Shot size / camera	Picture	What the shot gives the audience
S01	4 s	Medium shot, eye level, slightly behind the side	End of the corridor; the rider walks to the door with the bag; rain drips from the coat	Person, place, weather: establish the space
S02	5 s	Close shot, low angle, side-on	Crouches to set the bag down, straightens it, raises a hand and stops at the door	Action and hesitatio n; the hand is the point
S03	4 s	Close-up, eye level	Warm light in the door gap; the hand withdraws; he stands and leaves frame	The choice not to knock

Shot	Duration	Shot size / camera	Picture	What the shot gives the audience
S04	3 s	Close-up, low angle	Empty corridor and the bag; a cat paw reaches out, hooks the opening and drags it away	The turn and the aftertaste

Four shots are not arbitrary cuts. S01 sets up, S02 performs the action, S03 delivers the choice, S04 reverses it. Each advances one thing only, so each can be redone on its own without touching the film. That is the advantage of a short piece: if it is wrong, redo one shot, not the whole film.

The angles have reasons too: S02 sits low because the audience must see the hand stop at the door; S04 also sits low so the cat paw enters from under the door gap, echoing the eye height of S02. Four eye-level medium shots would thin the action out into a checklist.

Step 5: prompts - translate the picture into executable parameters

A prompt is not an essay; it states what is already decided, in a fixed order. Each prompt should have six blocks: subject and look, action, environment and light, camera and framing, texture and style, and exclusions. S02 is the example.

> Rainy night, corridor of an old residential block. A male delivery rider in his mid-20s, short black hair, dark grey raincoat with a reflective strip across the chest, black rain boots.
> He crouches in front of the steel door, sets the kraft paper bag on the ground and straightens it, then raises his right hand and stops his knuckles a few centimetres short of the door. He does not knock.
> A single cold white lamp at the end of the corridor lights him from behind right; warm yellow light spills from the door gap; fine droplets sit on the raincoat; shallow puddles on the ground catch the light.
> Close shot, low angle, side-on framing, 35 mm view, the figure occupies two-thirds of the frame on the right, the door panel on the left, camera locked off.
> Realistic film texture, restrained saturation, cold blue against warm yellow, shallow depth of field.
> No dialogue, no subtitles, no music, no extra people, no exaggerated expression, no deformed fingers, and the hand must not actually touch the door.

Every line of this prompt maps to a decision already made: the look comes from the character card, the action from the script, the light from the location card, the camera from the shot list, and the exclusions from known failures. Nothing says "crank up the cinematic feel", which cannot be executed.

Three forms that are most often missed: first, **write the negative** - "the hand must not touch the door" constrains the model better than "the hand stops at the door"; second, **write the count** - "one cold white lamp" is steadier than "a dim corridor"; third, **write the relative position** - "door on the left, man on the right" is more useful than "nice composition". AI does not understand taste; it understands relations it can check.

Step 6: generate stills, generate video, adopt

1. **Start with the first frame:** One first frame per shot; check the hand, the door gap and the bag. If the frame is wrong, rewrite the prompt and regenerate; do not take a bad frame into video.

2. **Only then generate video:** One main action per shot, duration from the shot list. Check model, resolution and duration before generating; points are charged only after you confirm.

3. **Review candidates, then adopt:** Generation only produces candidates; adopting is a separate action. Only adopted footage enters the timeline.

4. **Redo one shot, not the film:** Fix the shot that is wrong and keep the shots that passed. Do not re-run the whole film.

> **Cost boundary:** First frames, video and voice are priced separately, and a redo is priced the same. Confirm all four first frames before generating video in one pass; that is cheaper than generating and reviewing one shot at a time. Points come from the quote page; this tutorial does not quote prices.

Step 7: edit and sound - three things only for 16 seconds

The edit does three things only: hold long enough on the "hand stops" in S02, cut the "hand withdraws" in S03 on the frame the action completes, and bring the cat paw in S04 half a second earlier. No transitions, no filters, no music.

Sound comes in three layers: rain as the bed throughout, footsteps and raincoat friction in S01-S02, bag rustle and the soft door-gap sound in S03-S04. With no sound library, write the sound sheet first: which layer enters at which second and leaves at which second, then decide whether you need voice or generation.

Step 8: acceptance - let the file speak

□ About 16 seconds, four shots in the right order, no repeated action.

□ The hand in S02 really stops at the door, with no clipping and no extra fingers.

□ The hand withdrawing in S03 and the cat paw in S04 connect logically.

□ No dialogue and no subtitles anywhere; the rain is continuous and never cuts out.

□ Play the downloaded file from the start; picture and sound are both normal.

What this drill teaches

It teaches you **even a short piece takes the whole path**: a story has a turn, a script has actions, assets have cards, the shot list assigns tasks, prompts carry constraints, and the finished film has acceptance. Get a 16-second film right and 60 seconds is only more shots; the method does not change.

It also shows AI's real limits: hands, contact points, light through a gap, a paw gripping. Those are where models fail most. That is why prompts need negatives and relative positions, why first frames come first, and why a redo touches one shot. That experience is worth more than any prompt template.

Next: use Little Boat Across the Puddle to see another approach that tells a story with action and environment alone.

Drill 2 • Little Boat Across the Puddle: no characters, but still a protagonist

15 seconds, three shots, no people, no dialogue. The protagonist is a folded paper boat. It completes a departure, an obstruction, and a landing, so the audience sees a journey in one puddle.



If it will not play, download the 480p film (854×480, 15.000 s, 24 fps, 1.4 MB).

Finished film: 15 seconds / 3 shots / 854×480 / 24 fps / no dialogue. Three 5-second shots: a raindrop lands, the boat sets off, and it lands by a grass blade.

Step 1: story - write the object as a character

One-line story: before the rain stops, a folded paper boat sets out from a drip under the eaves, is pushed across the puddle by ripples, hits a fallen leaf and spins, and finally stops beside a dew-covered grass blade.

No people, but the story structure loses nothing: the boat has **starting point** (under the eaves), **driving force** (drips and ripples), **Obstacle** (a fallen leaf), **ending point** (the grass blade). The audience stays because it feels like a small voyage, not three pretty shots of water.

Writing an object as a character means giving it intent. The boat cannot row, but it can be pushed, blocked, and held. Ask three questions: Where does it come from? What pushes it? What stops it? Answer them and you have a story.

Step 2: script - no lines, only physics

Little Boat Across the Puddle, about 15 seconds, 16:9, realistic macro look, after rain.
1. Before the rain stops, the last drip falls from the eaves into the shallow puddle at the door, sending out rings of ripples. A folded paper boat floats on the water.
2. Ripples push the boat forward, its bow rocking gently as it crosses the middle of the puddle; the water reflects a grey-white sky and a little of the eaves' shadow.

3. A soaked fallen leaf lies across the water; the boat hits it, spins, and slowly stops beside a dew-covered grass blade, no longer moving.
No dialogue, no subtitles, no score, only the drip, the ripples, and distant birds.

This script has no emotional words; it is all shootable physics: drip, ripple, sway, collision, spin, stop. Sound is also three kinds only: drip, ripple, birdsong. The feeling comes from "it has finally stopped", not from a voice-over saying "this moment is quiet".

Step 3: character, location, prop - objects get three cards too

Three cards even with no people

Asset	What to write	Why
Protagonist · folded paper boat	Folded from white copy paper, creases in the hull, the bow slightly raised, pale grey water stains, about 8 cm long	Scale sets the macro ratio; the creases say hand-folded, not a plastic toy
Location · puddle at the door	Concrete ground, a shallow puddle about 60 cm wide, grey-white sky reflected in the water, moss and fine sand at the edge	It fixes what is reflected and the colour of the water, so no shot invents a new ground
Prop · fallen leaf and grass blade	One soaked brown fallen leaf with curled edges; a clump of fine dewy grass, blades 2 cm below the water surface	The leaf is the obstruction and the grass is the destination; both must be something the boat can touch

An object card must state **scale**. The same "folded paper boat" at 8 cm and at 30 cm gives two completely different images: macro and toy. State the scale clearly so the model knows how close the camera should be and how shallow the depth of field.

Step 4: shot list - three shots make a complete arc

Shot list (15 seconds / 3 shots, 5 seconds each)

Shot	Shot size / camera	Picture	What the shot gives the audience
s0	Close-up, low angle, just above the water	A drip falls from the eaves, ripples spread, the boat sways on the water	Starting point and driving force
s1	Medium close shot, side-on tracking	The boat is pushed forward by the ripples across the middle of the puddle	Journey and speed
s2	Close-up, low angle, locked off	It hits the leaf, spins, and stops by the grass blade	Obstruction and destination

Five seconds each is not averaging: each action needs time to be read. The drip needs the ripple to spread, the journey needs a change of speed, the stop needs the sway to settle. Under four seconds feels cut off; over seven seconds adds no information.

Every camera sits just above the water because that is the boat's eye height. From a person's height looking down, the boat becomes a small white dot in the puddle and the journey disappears. The low angle is not just prettier; it is necessary for this story.

Step 5: prompts - macro, reflection, contact point

> Shallow puddle at the door after rain, macro view, a folded paper boat about 8 cm long, creases in the paper and pale grey water marks.
> The boat hits a soaked brown fallen leaf, the bow turns aside, it turns slowly on the water, and finally stops beside fine dew-covered grass blades as the ripples settle.
> Grey-white sky reflected in the water, the eaves' shadow above, grass blades below the waterline, dew on the tips, fine dust floating on the surface.
> Close-up, low camera just above the water, locked off, shallow depth of field, focus where the boat meets the leaf.
> Realistic macro photography look, soft overcast light, restrained color, real water and paper texture.
> No dialogue, no subtitles, no score, no people, no exaggerated splash, do not deform or sink the boat, do not let the leaf drift away.

This prompt turns on three words: **macro** (sets the scale), **reflection** (sets what the water reflects) and **contact point** (sets the focus). The danger in macro work is that it looks like a big scene, so put the scale in: 8 cm, grass 2 cm below the surface.

Negatives matter just as much: do not let the boat sink, do not let the leaf drift away. Models love to add drama, and the strength of this small story is restraint-it stops, it is not wrecked.

Step 6: generate stills, generate video, adopt

1. **Three first frames first:** Check that the water reflection stays continuous, the paper creases match, and the grass and leaf keep their relative positions.

2. **Then three video clips:** One action each: drip, travel, stop. More actions per clip means more mush.

3. **Adopt, then go to the timeline:** Candidates are not the film; only adopted footage appears in the edit.

4. **Redo only the bad clip:** If s1 is too fast, redo only s1 and keep s0 and s2.

Step 7: edit and sound - leave only breath between shots

Cut the three shots straight, with no transitions. Let the ripple in S01 spread until half a second before the cut, join the bow entering frame in S02 to the edge of the ripple, and cut S03 on the frame the bow touches the leaf. Every cut lands on an action, so nothing feels like a jump.

Three sound layers: the drip (start of S01), water and light friction (S02), birdsong and the water settling (S03). Birdsong is not decoration; it tells the audience the rain has stopped, and with the grey-white sky it completes the time of day.

Step 8: acceptance - did it finish the journey?

□ About 15 seconds, three shots in the right order, and the boat moves in only one direction.

□ The water reflection is continuous; no sunny shot followed by an overcast one.

□ The boat stays the same size; it never suddenly grows or shrinks.

□ The leaf and grass keep a believable relation, and the collision is not clipping.

□ No dialogue and no subtitles; the drip, water and birdsong layers are continuous.

□ The downloaded file plays through and ends at the grass blade, not on a black frame.

The two drills side by side

One process, two kinds of difficulty

Stage	Don't Knock	Little Boat Across the Puddle
Story	The turn comes from an action that does not happen	The turn comes from an object's journey
Character	Character look and hand action	Object scale and material
Set	Corridor space and light in the door gap	Water reflection and ambient colour
Prop	Bag opening and handles	Leaf as obstruction, grass as destination
Shot list	Four shots, one piece of information each	Three shots, one complete action arc
Prompt difficulty	Negatives and contact points	Macro, reflection, scale
Acceptance focus	Hands do not clip; the paw can grip	Reflection is continuous; size stays the same

Both films are small, yet together they cover the two problems bigger films also meet: **action consistency with a character** and **physical credibility without one**. After practising both, the

five full-case films become easier to read, because you can see why each shot is built the way it is.

Next: go back to Seven stages to check the output of each stage, or jump into Case 1 and make your first film.

From story to finished film: seven stages, and what AI and prompts handle

Don't Knock and Little Boat Across the Puddle follow the same chain. What matters is not the seven names but what each stage leaves behind, who makes the judgement, how far AI can take over, and what a prompt must actually constrain at that step. Skip a sentence here and you pay for it with a redo later.

Seven stages: every step must leave something checkable

Stage	Output	Human decision	AI can take over	What the prompt must pin down
1 Story	One sentence: who wants what, what blocks them, what they finally choose	What the story is and where the turn sits	Offer several directions, find causal gaps	Write character - desire - obstacle - choice - result, not mood adjectives
2 Script	Scenes, actions, dialogue or narration, opening and ending	Tone, pace, what is left unsaid	Expand into filmable actions, catch abstractions	Turn "sets it down gently" into "crouches, straightens the bag, stops his hand at the door"
3 Character	Character sheet: age, build, hair, clothing, which hand holds what, identifying details	Which feature must not change across shots	Extract traits from the script, produce a character sheet and reference images	Write checkable appearance, not "handsome" or "full of story"
4 Scene	Scene sheet: layout, entrances, light sources, materials, time and weather	Spatial relations and light direction	Generate wide views, details and lighting references	State where the door, counter, lamp and character stand relative to each other
5 Prop	Prop sheet: size, material, wear, text, contact points	What narrative job the prop does	Add material detail, generate front and detail references	Make "an old keychain" brass-coloured, scratched, engraved on the back and a specific size

Stage	Output	Human decision	AI can take over	What the prompt must pin down
6 Storyboard	Shot list: number, shot size, camera position, lens, duration, action, sound	What the audience learns in each shot	Break the script into shots, flag repeated information and gaps	One main action per shot, with a clear start, end and reason for the camera move
7 Finished film	First frames, clips, ambience, edit, subtitles, exported file	Which take is accepted, whether the rhythm works, whether it can ship	Images, clips, frame interpolation, mixing suggestions	Video prompts constrain action, camera, light, duration and prohibitions together

The real line is not whether AI can write but whether judgement can be outsourced. What the story says, what the character must look like, which take is accepted and whether the film can ship are human calls. AI is good at expanding decided things into more options and at making repetitive work fast. A prompt is not a spell; it is an executable agreement that pins appearance, action, space, light, camera and prohibitions down to the point where they can be checked.

Prompts change shape at every stage

One story, seven stages, seven different prompts

Stage	What the prompt has to solve	What goes wrong
Story / script	Constrain structure, motivation and filmable action	It reads as mood prose and cannot be broken into shots
Character / scene / prop	Constrain identifying details, scale, material and references	A new face, a new store or a new prop in every shot
Shot list	Constrain shot size, camera, lens, duration and order of information	Everything is a medium shot and the audience does not know where to look
Film	Constrain action, camera move, light, contact points and negatives	The model adds drama, hands intersect, props deform, timing does not match

Where each walkthrough breaks

The hard part of Don't Knock is **negative actions and contact points**: the hand must stop at the door instead of knocking, and the cat's paw must hook the bag instead of passing through it. The prompt has to say "does not knock" and "the hand must not touch the door", and pin where the hand sits relative to the door and the paw relative to the bag.

The hard part of Little Boat Across the Puddle is **scale and reflection**: an 8-centimetre paper boat, grass blades 2 centimetres below the water line, and a continuous sky reflection. The prompt has to fix the macro scale, what the water reflects and where the contact happens, or the model turns a puddle into a lake and the paper boat into a plastic toy.

Together they make one point: a prompt is not better for being longer, it is better for being checkable. Every stage leaves a product that can be inspected, so the next stage does not have to guess.

Acceptance checklist

□ You can say what each of the seven stages leaves behind, not just which buttons you pressed.

□ You know which calls a person must make and which execution can go to AI.

□ You can name the stage where your last failure happened instead of saying "the model is bad".

Review: Rewatch any finished film and find one place where the script has something the frame does not show, or the frame shows something the script never wrote. That is usually the break.

Next: use Advanced creation to split these judgements further, or go straight to Eight lessons and start making.

03 Eight Starter Lessons
the shortest manual path

Get to know the creation space

Lesson 1 of 8

First know where you are, and where your content is saved.

> **Lesson output:** A free canvas with clear ownership.

Location in TimeGleam

> Home → Canvas creation → New blank canvas / More ways to create; inside a work, top bar → Mode and canvas

1. **Confirm the workspace:** The workspace determines member permissions and available credits. Admins and editors can create; view-only members can browse and open existing content but cannot generate or modify it.

2. **Choose the right entry point:** The desktop home page has two columns, "Canvas creation" and "Workflow creation", stacked on narrow screens. "New blank canvas" carries no current brief or attachments; "More ways to create" is collapsed by default, and when expanded you choose text, image, video, audio, or "Start from materials", which brings in the applicable brief and materials and only creates a draft work, EP01, and a free canvas - it submits no model task. These eight lessons use this manual route first.

3. **Tell automatic from manual:** "One-click canvas plan" calls the model to build an editable plan, and it is also billable; "One-click canvas film" plans and prices first, then starts continuous production only after explicit confirmation. The workflow column also has New work, Import a plan, and Quick work. Don't use the auto-generate buttons in place of "enter draft".

4. **Find the original work:** "Works in this space" at the bottom of the home page continues content in the current workspace; "My works" at the top left covers all accessible workspaces, and "My assets" covers materials. Check the work, episode, and specific canvas before opening. " Mode and canvas" at the top separates the fixed workflow view from independent free canvases; a fixed view cannot be deleted like a normal canvas.



Check the space / work / episode first; workflow mode and the workflow canvas are fixed views, and a named free canvas is edited separately. An episode shares work context while free documents are saved separately. Interface check: 2026-10-16, example data.

Completion criteria

□ You can find the current workspace, work, episode, and canvas name.

□ You know a free canvas doesn't require a formal script first.

> **Common mistakes:** Creating new works over and over just to find a node. Look in "Mode and canvas" of the original work and episode first.

Start with an idea

Lesson 2 of 8

Turn one line of an idea into a clear, shootable brief.

> **Lesson output:** A short film concept and a written draft of 5-6 shots.

Location in TimeGleam

> Home canvas → Enter brief → More ways to create → Text; canvas left side → Text / AI assistant

1. **Define the boundaries：** Specify the characters, scene, action, style and target duration. Choosing Text on the home page carries the brief into a text node and does not start generation immediately.

2. **Have the assistant organize it：** Open the creation assistant from the left-side AI assistant or Ctrl+/, reference the current node, and ask for a plan first. Check the references in this turn so the assistant doesn't guess at other canvases.

3. **Save usable text：** A regular text node is not the same as a formal script. AI candidates must be explicitly replaced or appended; review the differences before accepting, so the result isn't only discussed in chat.

4. **Prepare the next step：** One main action per shot, keeping the same character, scene and prop. Make S04 first and confirm the look before expanding.



Canvas and workflow entries are separate; More creation is collapsed by default and the eight lessons start with a manual draft. Works in this space keeps concrete creation entries. 2026-10-19, current components with example data, not a daily business space.



Expanding shows manual entries only and calls no model; New blank canvas carries no brief or attachment, while the other entries carry content by purpose. 2026-10-19, interface illustration.

> Draft a short film plan of about 30 seconds at 16:9 for "The Last Lamp": an old grocery store on a rainy night, where 25-year-old Lin Mo finds an old brass keychain engraved with "Go home" and takes it after hesitating. Use only 1 person, 1 location and 1 prop, split into 6 shootable shots. Give the text plan first, without submitting image or video generation; state the action, shot size, sound and suggested pacing.

Completion criteria

□ Text is saved and can be viewed again.

□ Every shot has a clear action, with no characters or scenes added unintentionally.

> **Common mistakes:** Treating "the plan is written" as if the assets were already generated. Judge completion by the save receipt and task status.

Prepare assets and subjects

Lesson 3 of 8

Organize the looks you'll reuse, then reference them in shots.

> **Lesson output:** Reusable references for Lin Mo, the shop and the keychain.

Location in TimeGleam

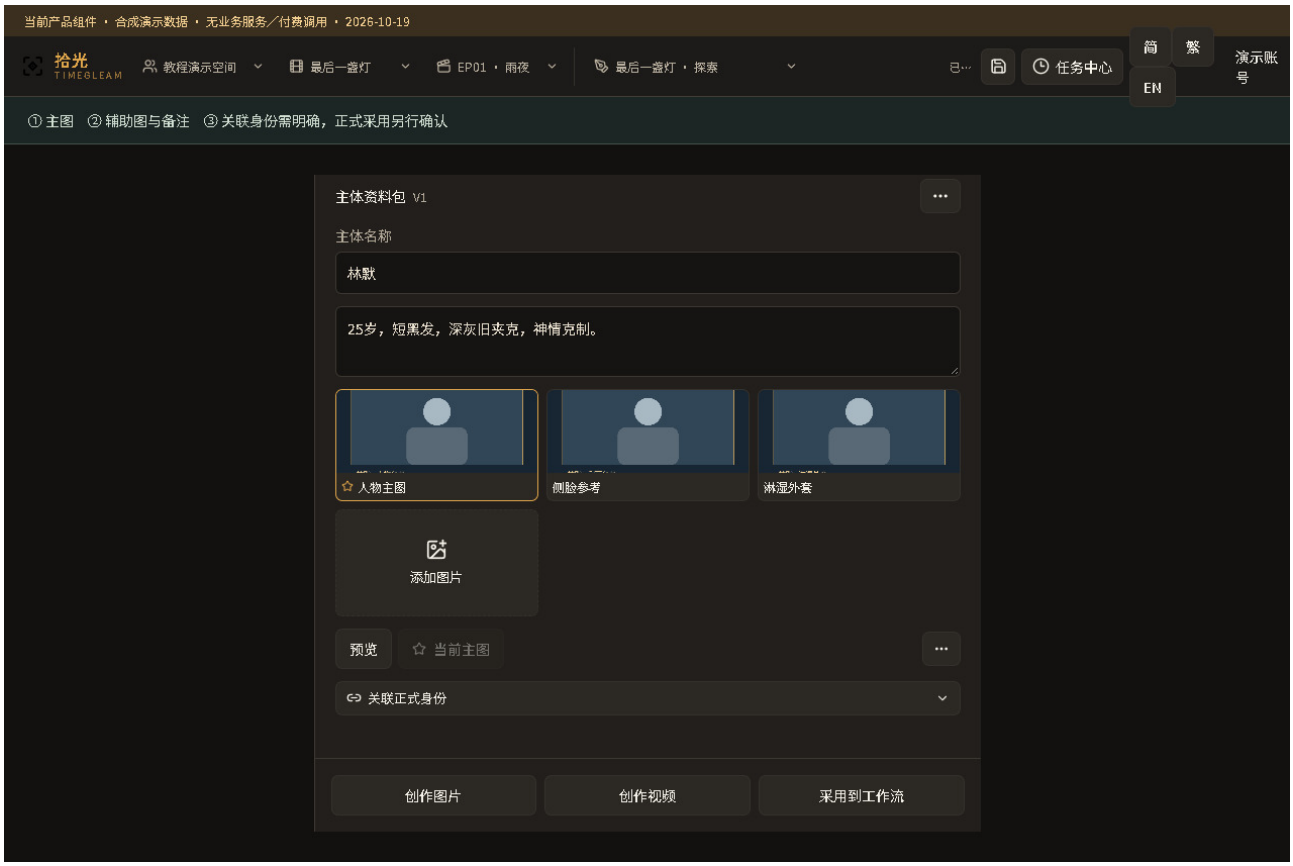
> Left side → Asset center / Upload; Add menu → Subject; creation description → @

1. **Tell the three object types apart:** Media assets are files; official characters / scenes / props are identities and versions; subjects are free-form creative reference packs. Changing only the media category does not create an official identity.

2. **Organize subjects:** The Add menu's "Subject" creates a new pack where you can set a primary image, supporting images, notes and order, up to 12 images. Removing one doesn't delete the file; find existing subjects under the Subject category in the asset center.

3. **Make a real reference:** Use @ to select, click a reference strip to insert, or drag assets into supported nodes. Uploading in the description field adds directly to the current node; tags can be clicked to preview or removed as a whole.

4. **Check the snapshot:** A subject reference saves a snapshot of that moment. Editing the original pack won't update existing references automatically; removing a reference doesn't delete the source file, so check the scope before referencing again.



The main image fixes the overall look and supporting images add expression and styling; sound and formal identity are set as needed. Linking is not formal adoption, and supporting images and sound are not written into a formal character version automatically. Interface check: 2026-10-16, example data.

Lin Mo: 25, short black hair, restrained expression, worn dark grey jacket, black trousers.
Scene: old grocery store on a rainy-night street corner, wooden counter, warm yellow light, rain streaks on the glass window, wet deep-blue street outside the door.
Props: worn oval old brass keyring, "回家" engraved on the back, scratched edges. Use clear, easy-to-read backgrounds for the character and prop references.

When do you need formal visual assets?

When you need formal shot binding, cross-episode versions, and batch management, create versions per Workflow visual assets subject. A subject can link a defined identity and version in this work; adopting into the workflow requires preview confirmation, a formal version adopts only one main reference image, and existing shot bindings do not upgrade automatically.

Completion criteria

□ The target node shows the actual reference.

□ You can tell source files, asset packs, and formal visual versions apart.

Common mistakes: Saying "reference the previous image" in the text only, without an actual reference. Check the thumbnails and the list before submitting.

Make the first frame

Lesson 4 of 8

First, lock the visual direction of one shot.

> **Lesson output:** One usable frame of S04, or a clear visual reference for it.

Location in TimeGleam

> Images on the left, or double-click / right-click empty space → Image; result → Preview / Edit

1. **Describe the frame:** Write the character, the moment of action, the scene, shot size, and lighting; actually add character and scene references; avoid contradictory requirements.

2. **Choose model settings:** Use a model available in the current workspace. Aspect ratio and size follow the model and the use case; not every image model supports a separate ratio; the chain button means follow the default.

3. **Compare candidates:** After submitting, watch the original task; when it finishes, review the candidates and adopt one explicitly. Check identity, costume, props, and composition first, then aesthetics.

4. **Edits you need:** Preview / Edit lets you crop, annotate, rotate, or mirror; after previewing, save as a copy and keep the original asset. AI layer separation and regeneration are billed per actual call.



The description says what to shoot; thumbnail chips show the actual references, click to preview and remove the whole chip; More settings shows input details. A material name in text alone does not guarantee the model receives the image. Interface check: 2026-10-16, example data.

> Lin Mo stands at the counter of the old grocery store, eyes fixed on the old brass keyring, leaning slightly forward. Close shot, warm yellow counter light on his side face, deep-blue rainy night beyond the window. The dark grey jacket and short black hair match the character reference. The beat is "discovery"; do not show him picking it up or leaving yet.

> A composition preview is not a required step for every video shot. With existing references you can go straight to video; a strict first-frame / first-and-last-frame route still needs the matching images.

Completion criteria

□ Character, scene, and props match the story.

□ You can name the one key action of this shot.

> **Common mistakes:** One image looks great, but the costume keeps changing. Go back to the subject and references, fix it, then do the next shot.

Bring the FRAME to life

Lesson 5 of 8

Leave the static look to references; put action and camera movement in the prompt.

Lesson output: One usable clip of S04.

Location in TimeGleam

Image result → Generate video; video node → description, references, model, specs, generate

1. **Create a draft:** "Generate video" on an image result creates a video draft connected to the original image; you still need to add a prompt and submit-this step does not generate immediately.

2. **Check the reference mode:** Choose multimodal, first-frame, or first-and-last-frame mode as supported by the model. In strict first-frame mode, the aspect ratio follows the first frame; if old specs are incompatible after switching models, select again.

3. **Choose valid specs:** Use duration and resolution options supported by the model. Higher resolution or longer duration may cost more credits; selected specs do not equal the actual output pixels.

4. **Submit one shot:** Check the description and the image, video, and audio reference list. After submitting, monitor the original task; do not regenerate repeatedly while waiting for the result.

Lin Mo's gaze slowly falls to the old brass keychain; his body leans slightly forward, and at the end of the shot he reaches for the keychain. The camera pushes in slowly, warm yellow light stays steady, rain runs down outside the window. Keep the character's appearance and the space continuous; do not cut to another scene.

Handle sound when there is dialogue

Submit audio references for regular video only when the model supports them. In a production workflow, characters with dialogue need their character voice completed-preview it first; a locked look does not mean the voice is ready. Audio in the asset pack does not automatically become a workflow character voice.

Completion criteria

□ You have watched the actual video, not just seen a successful submission.

□ The action, camera movement, and references meet the requirements.

Common mistakes: A single shot asks for picking up, flipping through, paying, turning off the light, and leaving. Keep one action first; split the others into the next shot.

Finish and fix the other shots

Lesson 6 of 8

Keep the results that work first, then fix the most obvious problems.

Lesson output: 5-6 clips ready to edit, with broadly consistent characters and style.

Location in TimeGleam

Assets and versions → candidates; selected node → send to assistant

1. **Expand from the plan：** Reuse references and write a separate action for each shot. Duplicates distinguish between keeping input references and not including input references.

2. **Compare and use：** Assets and versions show "Latest generation · Not used" or "Currently in use". Viewing does not change the version; candidates must be explicitly used.

3. **Local edits：** Select the node and send it to the assistant, specifying that only this shot should change. Review the references, differences, and scope before using it; saving edits is not the same as regenerating.

4. **Keep old results：** Old results pending update remain viewable. Using a new result does not guarantee automatic replacement on the timeline; check which version is used when editing.



Check the current adoption marker first, then preview candidates and adopt explicitly. Generating a candidate, adopting in the node, adding to the timeline and adopting into the workflow are four different actions. Interface check: 2026-10-16, example data.

> Work only on the S04 I reference: keep the character, scene, costume and existing references, and change the action to "his gaze falls to the keychain and he leans forward slightly", without adding a pick-up or an exit. Give me the differences first and do not redo the other shots. When generation is needed, list this round's targets for me to check.

Six-shot rhythm example

1. A rainy night street corner; the last light in the small shop is still on.

2. Lin Mo pushes the door open and closes the umbrella; the doorbell rings.

3. Finds an old brass keychain.

4. Close-up: turns it over to the back, which reads "home".

5. Cash left on the counter; the lights go out.

6. He walks into the night, the keychain in his palm.

Split by generatable actions, aim for around 30 seconds in the edit, and don't force a single asset to equal the suggested shot length.

Completion criteria

□ Each shot has a usable result, and the character's costume stays consistent.

□ Know which version is currently in use, and keep usable history.

Common mistakes: Redo all upstream work just to fix one hand. When generating only the selected targets, first confirm the upstream has valid results you can reuse.

Editing and sound

Lesson 7 of 8

Arrange usable shots into a film with a beginning and an end.

> **Lesson output:** A complete first cut that plays from start to finish.

Location in TimeGleam

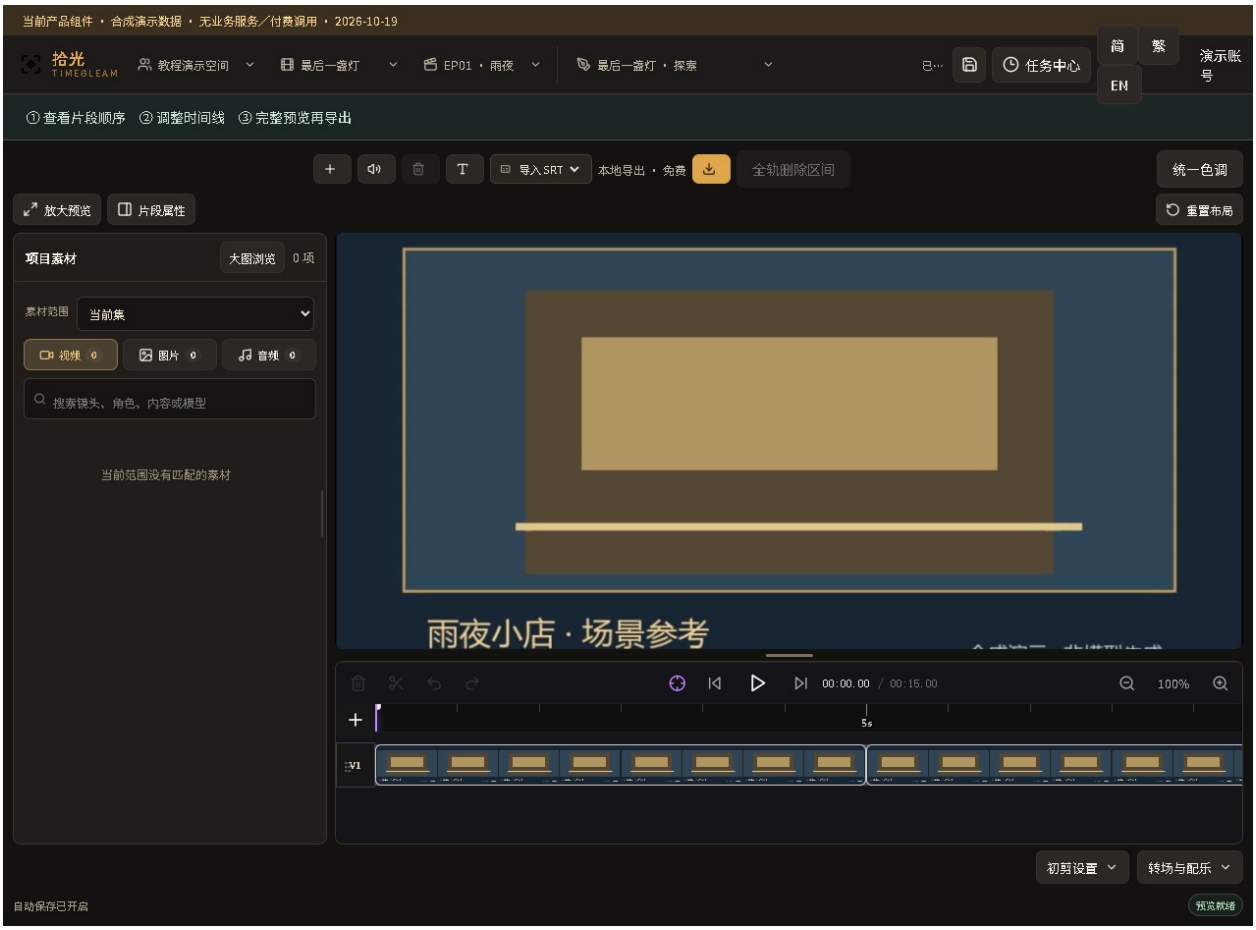
> Result → More asset actions → Add to this canvas timeline; top → Timeline

1. **Place clips in order:** Add the video to this canvas timeline. In workflow canvases, the equivalent entry point is Add to continuous preview; for full editing, see the dedicated topic.

2. **Adjust the pacing:** Adjust the order and in/out points, keeping the usable action. Play it all the way through before fixing slow or jumpy sections.

3. **Check sound and subtitles:** Listen to the original audio, dialogue, and score so they don't mask one another, and align subtitles with the picture. Use a computer for fine multi-track work and mixing.

4. **Review source assets:** Check the clip's source asset and trim range. After the canvas adopts a new version, check whether the timeline still references old assets, then handle them explicitly.



Order carries the story; trimming keeps the action; preview before export. This image uses current components and locally assembled footage, not a real generated film. Interface check: 2026-10-16, example data.

Completion criteria

□ It plays through, and the story is understandable.

□ No obvious bad frames, sudden character changes, or volume problems.

> **Common mistakes:** Fine-tuning after only two seconds of playback. Note overall issues first, and prioritize fixes that affect story comprehension.

Export and continue creating

Lesson 8 of 8

Confirm completion with the actual file, and leave a clear entry point for the next round of changes.

Lesson output: A downloaded and checked short film file.

Location in TimeGleam

Current timeline → Export; task details / generation history

1. **Save and export:** Check the timeline and specs, then check the task after export. It is complete only when the task succeeds and the file is usable.

2. **Open the actual file:** Check the opening and ending, duration, aspect ratio, audio, and black or frozen frames. In a formal workflow, you can start the basic file check for this version on the delivery page.

3. **Save the entry for next time:** Note the work, EP, and canvas name. On the home page, Continue recent work restores the original document; check the history, tasks, and adoption status separately.

4. **Restore failed:** First read the original task's reason, then use the restore or failure branch entry to reuse successful results. Pausing subsequent steps does not mean canceling tasks already claimed by the provider.

Billing follows the rules of the current workspace: text reasoning, speech recognition/read-aloud, and model generation may be charged; ordinary local editing, references, and previews do not incur an extra per-action fee. This tutorial does not fix prices. For permission or balance issues, check under Account/Workspace; models and providers are managed by platform administrators.

Completion criteria

□ You have opened the actual file and watched it all the way through.

□ You know how to return to the original canvas, check the original task, and continue editing.

Common mistakes: Don't treat the assistant saying "Completed" or an export being queued as delivery. Go by a playable, downloadable file.

04 Finished-Film Cases
Scripts, shot lists and layered production for five full short films

Case 4: A complete production • from script to The City of Spare Charge

What you see first: A 75-second industrial science-fiction short with 13 shots, attached as a 720p master and a 480p platform export. Every frame is made inside this project: block out the 3D previz to lock the staging, the lighting and the camera move, then render it out frame by frame. The head and tail are text cards. Model-assisted refinement is an optional branch that this cut does not use. Watch the film and you will see which stages a short goes through from script to delivery, and what evidence each stage leaves behind.

Studio -> Canvas -> New blank canvas -> Subject cards and shot nodes -> Manual film -> Export

See the result first: two files, two purposes



If playback fails, download the 720p master (1280x720, 75.02 s, 24 fps).

The 75-second complete production of The City of Spare Charge. Thirteen shots: title card, highway, figure, tower interior, socket, plant, corridor, whip pan, vault, cell landing, recovery, night city, end card. Watch it twice: the first pass for the story, the second counting shots and noting every cut.



If playback fails, download the platform 480p export (854x480, 4.1 MB).

The same timeline through Manual film -> Export in the platform. Platform export is fixed at 854x480; the 720p master is re-encoded locally from the same edit and does not pretend to be a platform export. Acceptance checks the real resolution, duration and audio track of the platform file.

Production chain: seven stages, each with an artifact

From a one-sentence story to a playable short

Stage	What you do	Artifact and acceptance evidence
1 Script	Logline, theme, five-act causal chain, shootable scenes	The script section of the production bible; one sentence must answer who lacks what, who stops him, and whether he gets it
2 Shot list	Shot size, camera, lens, movement, key light, duration, source for every shot	A 13-shot table totalling 1800 frames and 75 seconds, within one second
3 Subject cards	One card each for character, set and prop, fixing proportions, materials and colour temperature	Three subject cards with their supporting images; no drift across shots
4 3D previz	Block the low-poly set, lock camera and movement, render in layers	1280x720 frame sequences, no missing frames per shot; character and set layers separate
5 Optional refinement	Only when you need a different texture, feed a previz frame as the first frame and generate video candidates	At least two candidates per shot, previewed in full before adoption; not used in this cut
6 Edit and sound	Order by the shot list, trim, place two black gaps, lay ambient and action sound	A master timeline whose cuts land on the wind-up of each action
7 Export	Export 480p from the platform, re-encode 720p locally in the same order	Both files stream from a direct link; metadata matches the table

Step one: write the script first, not the pictures

A script answers only three questions: what he lacks, who stops him, and whether he gets it. Compress the answers into one causal chain - **no power -> take power -> blocked -> take it back -> light up** - then give each act its time. These timecodes are the actual cuts of the finished film; rehearse the rhythm against them while writing so the structure never has to be rebuilt later.

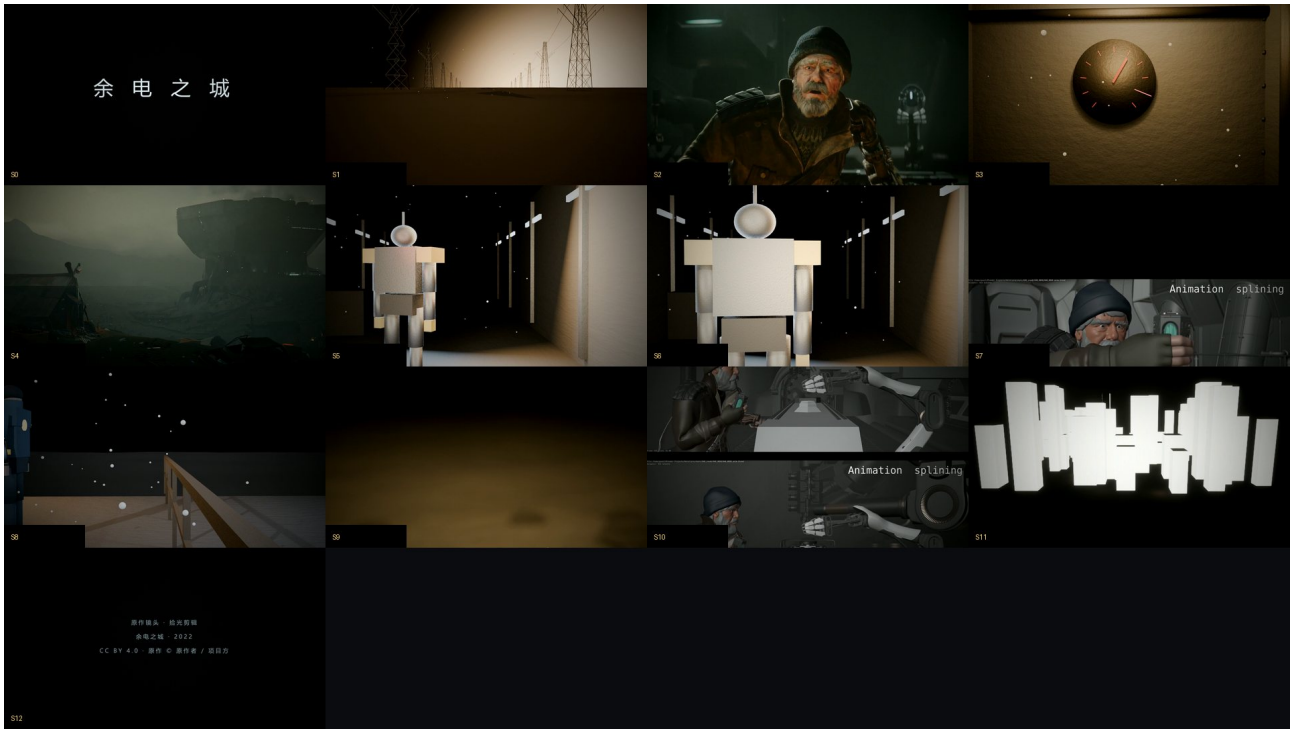
Five-act structure and causal chain (matches the 75-second film)

Act	Timecode	Function	Key action	emotion
0 Title	00:00-00:04	Set the tone, give the title	Title lights up character by character	Still
1 World	00:04-00:21	Establish that power is currency	Slow push down the highway; a figure's back faces the dead city	Bleak, curious

Act	Timecode	Function	Key action	emotion
2 Motive	00:21-00:33	The cell is dying; he must take power	The gauge in the tower drifts into the red; a hand pushes the cell into the socket	Urgent
3 Blocker	00:33-00:47	A patrol spots him	The robot turns; an over-the-shoulder follow; a whip pan across the plant	Pressure
4 Climax	00:47-01:04	Vault, fall, recover	Side-tracking vault; the cell hits the gravel; a dive pins it down	Peak
5 Resolution	01:04-01:15	The city lights up again	Window lights spread across the night city; closing card	Release

Do not describe how cool the camera is while writing the script. Let each act carry one action only, and move the cut three frames earlier than the wind-up so the audience never feels that the event finished before the cut.

Step two: the 13-shot list turns the story into executable parameters

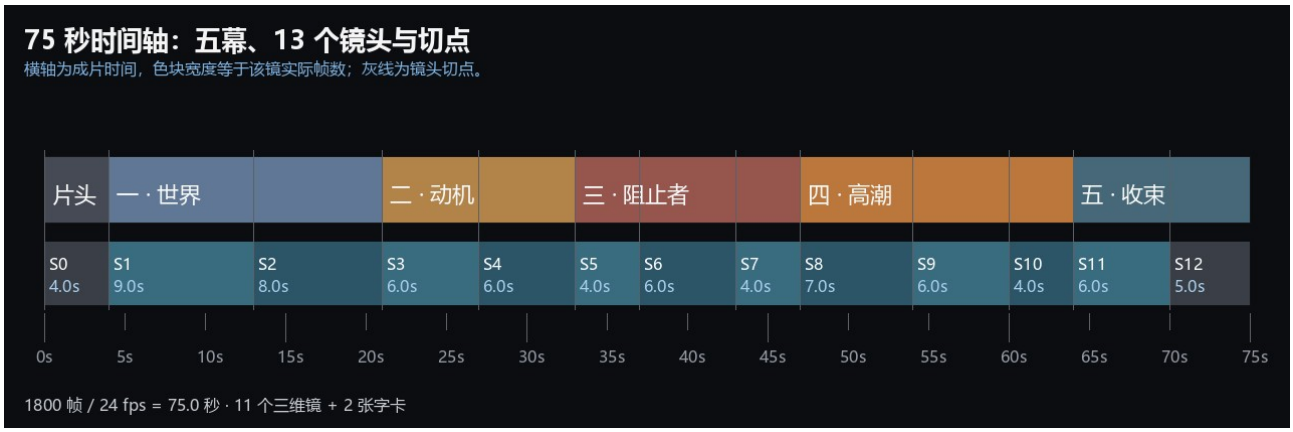


The 13 shots in film order. Read the rhythm of shot sizes first: extreme wides carry the world, medium close-ups carry the figure, close-ups are reserved for the cell and the hand; never use two identical sizes back to back inside one act.

Shot list for The City of Spare Charge, with actual render parameters

#	Shot size	Camera / lens	Movement	Key light and colour temperatur e	Duration	Source
S0	Title card	Front	Character-by-character reveal	Pure black, warm-white type	4 s	Title car d
S1	Extreme wide	Low angle 24 mm, 1.1 m above ground	Slow push 31 m, tilt 1.4 deg to 2.8 deg	Backlit sunset 3200 K, volumetric dust	9 s	3D prev iz
S2	Medium close	Eye level 45 to 52 mm, camera retreats 1.5 m	Slow push, light handheld	Side backlight 4300 K, cool fill	8 s	3D prev iz
S3	Medium	Eye level 35 mm	Lateral move 2.3 m	Warm practical 3000 K inside the tower	6 s	3D prev iz
S4	Close-up	High angle 85 to 76 mm	Micro push	Single hard light 3600 K	6 s	3D prev iz
S5	Wide	High angle 30 to 28 mm	Lateral move 2.6 m	Industrial overheads 5000 K, flicker	4 s	3D prev iz
S6	Medium	Over-the-shoulder 38 mm, 1.2 m behind the shoulder	Follow push 0.7 m	Warm front 2800 K / cool back 6500 K	6 s	3D prev iz
S7	Wide	Wide angle 24 to 28 mm	Handheld whip pan	Flickering plant lights	4 s	3D prev iz
S8	Medium wide	Side tracking 40 mm, 1.5 m above ground	Lateral tracking 0.9 m	Cool blue key 6500 K + warm rim 3000 K	7 s	3D prev iz
S9	Close-up	Low angle 85 mm, f/2.2	Slow push 0.18 m	Single amber highlight	6 s	3D prev iz
S10	Medium	Handheld 85 to 72 mm	Low-angle push 0.2 m	Warm flare 3200 K	4 s	3D prev iz
S11	Extreme wide	High angle 28 mm, 40 m above ground	Slow pull-back 35 m, descend 14 m	City lights 2700 K + night sky 8000 K	6 s	3D prev iz
S12	Title card	Front	Line-by-line fade-in	Black with warm white	5 s	Title car d

75 seconds and 1800 frames in total. The source column is not decoration: title cards, 3D previz and AI-generated shots must be labelled separately in final delivery so viewers and collaborators know how each frame was made.



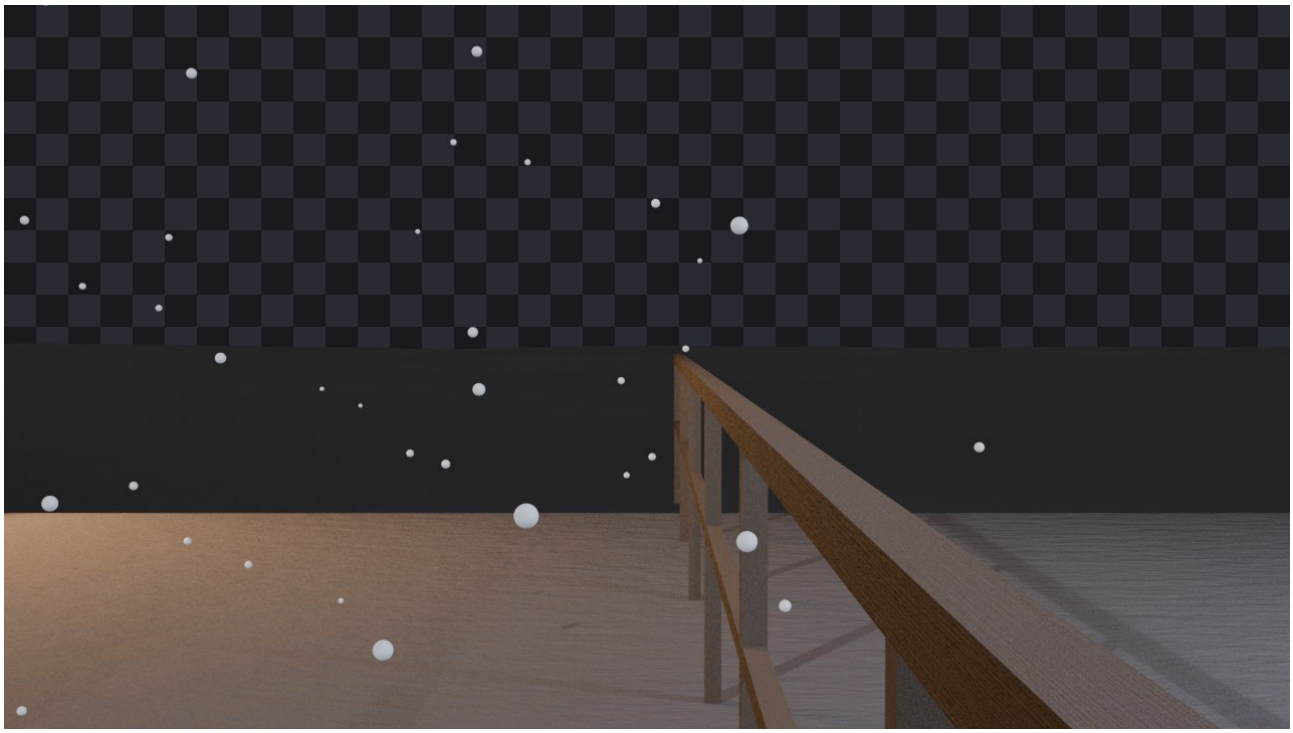
Draw the shot table as a timeline and length problems become obvious: S1 and S2 together run 17 seconds, nearly a quarter of the film, so the world has to land in one pass; the two close-ups S4 and S9 get 6 seconds each for the battery and the hand. Block widths come from real frame counts, not from a feeling. Change the length on this chart first, then change the shot table.

Step three: three subject cards for character, set and prop

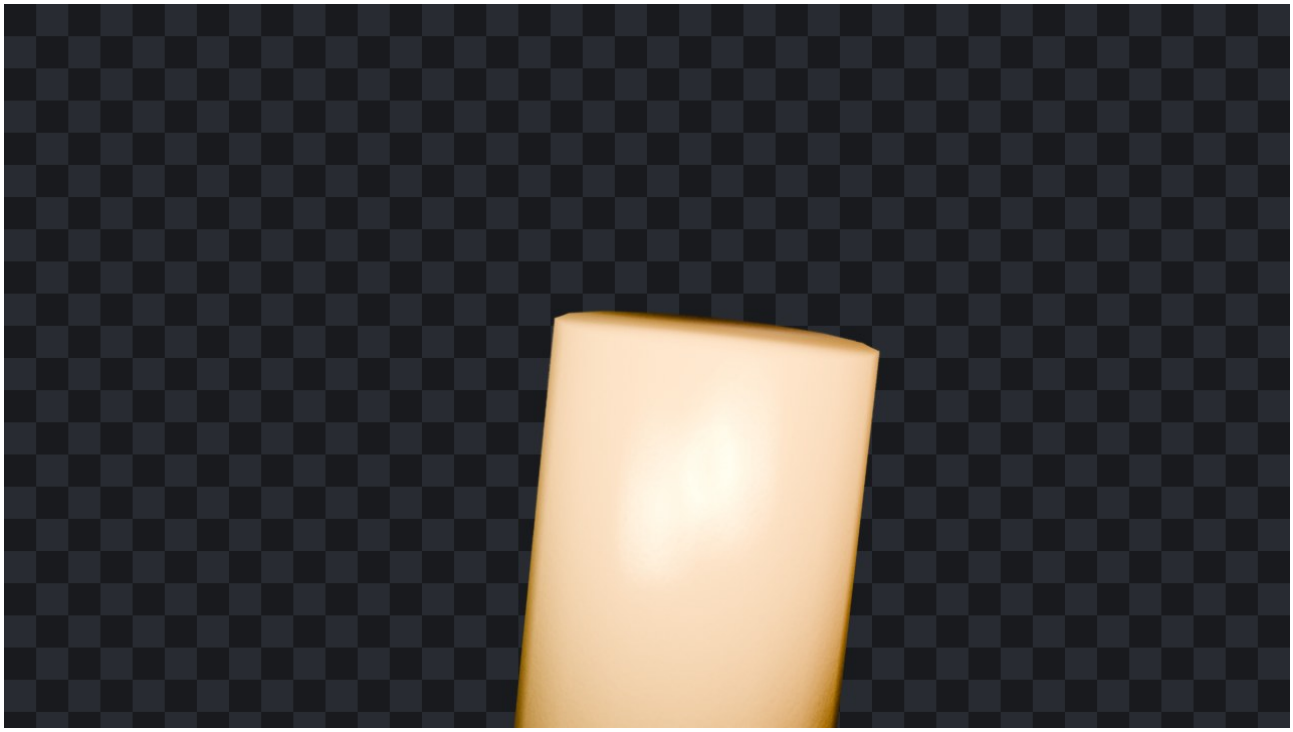
Fix who, where and what he carries before discussing action. Each card answers only appearance: proportions, materials, colour temperature, wear. Action belongs in the shot list; mix it into the subject card and every action revision quietly rewrites the look.



Character card: Shen Yan, a lean adult man in faded blue-grey workwear, bare metal on the left forearm with a band of warning yellow, an amber charge gauge on the left chest. Write anti-drift rules as negatives: the arm never switches sides, no cape, helmet or mask.



Set card: abandoned substation. State the scale (corridor width, clear height, depth), the materials (old concrete, steel, warning yellow) and the light (high-window side backlight plus warm practicals). A set card only records what does not change between shots.

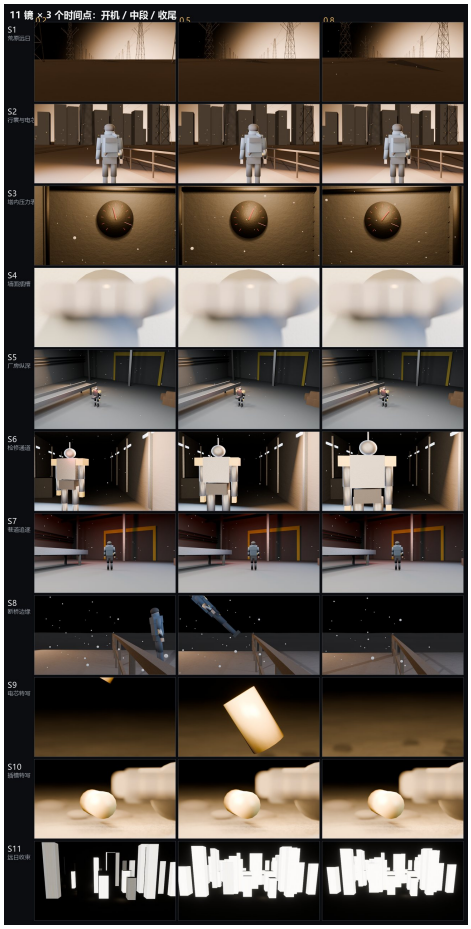


Prop card: storage cell, 24 cm long, one-handed, translucent amber shell, glowing core. A prop must say how big, of what material, and who holds it how; write only a battery and every shot invents a different object.

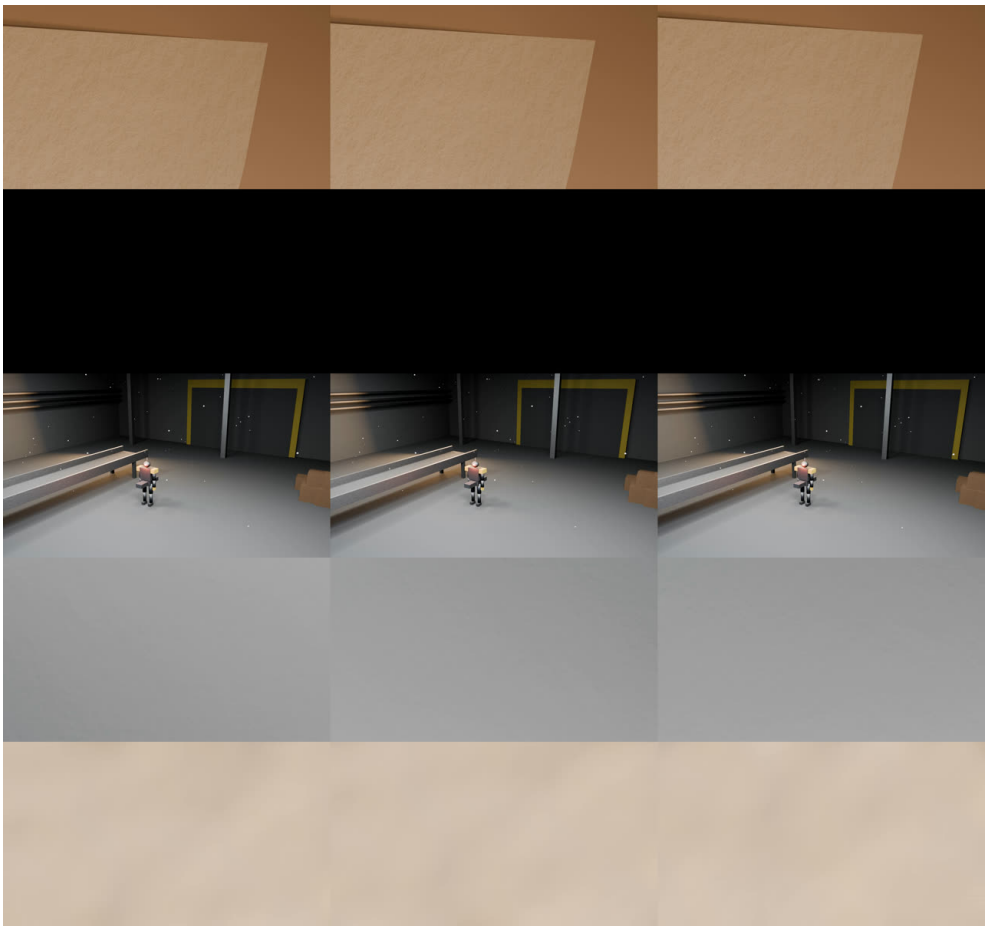
With the three cards in place, cross-shot consistency has checkable anchors: the back in S2, the vault in S8 and the dive in S10 must share one jacket, one mechanical arm and one gauge.

Step four: 3D previz locks blocking and light positions first

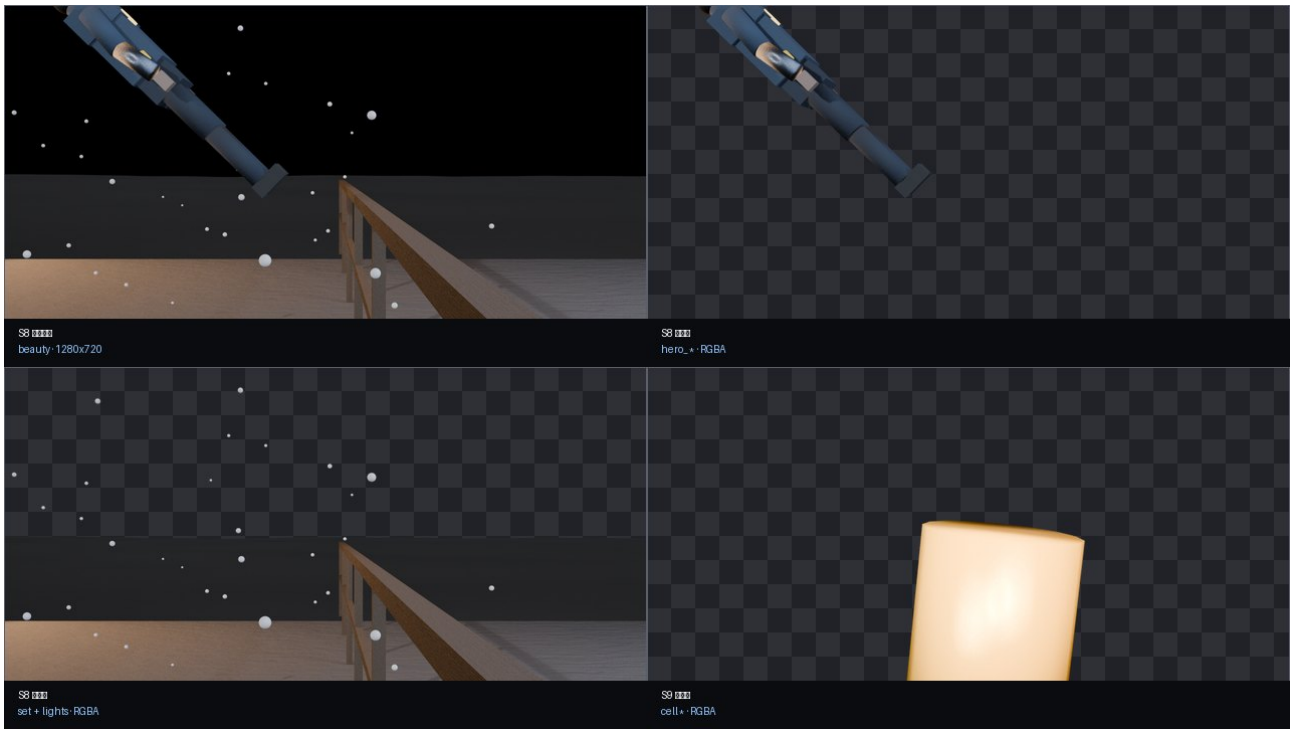
Previz is not "a quick animation for fun". It settles the three things generative video struggles to hold: **spatial relationships, camera movement, continuous light positions**. Block the low-poly set, lock the camera, the focal length, the motion curve and the key light direction, and only then decide whether any single shot needs model refinement.



11 3D shots x 3 time points (0.2 / 0.5 / 0.8). The three frames in each row must read as one continuous move: where the camera goes, where the subject goes, and whether the light drifts. This matrix is the first pass of per-shot acceptance: find the jump in the small images, then go back to the project and fix the keyframes; a fault in the middle of a move is easier to spot here than in the video.



The same five shots sampled at three moments (0.2 / 0.5 / 0.8 of each clip). One row is one shot: only the camera position, the focal length and the subject move, while the lighting and the exposure stay put. Read row by row to see whether the motion keeps a single direction. If the middle frame jumps, go back and fix the keyframes instead of patching it in post.



Layered render comparison: character and set rendered separately, then composited. The character layer can be relit and re-rendered alone without recomputing the whole plant for one shoulder line; the set layer can lock composition before the character is final.

Split character, set and prop into separate passes

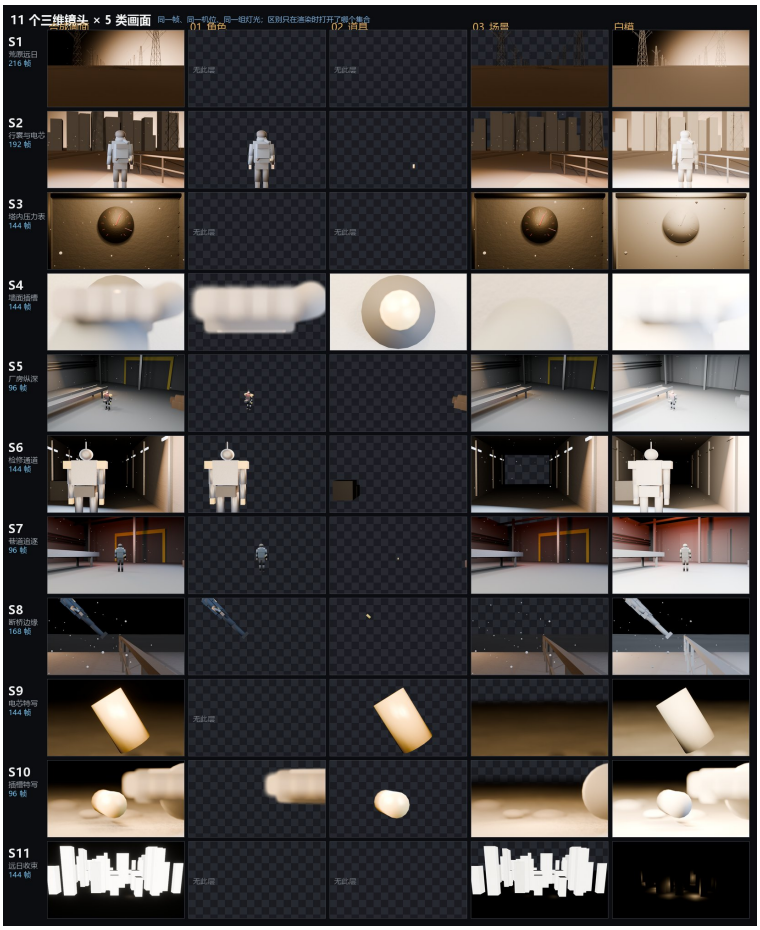
This section is the most valuable part of the case: instead of rendering the whole scene in one go, every shot is split into parts you can redo on their own. Each of the 11 3D shots is saved as its own project file, and inside it every element is sorted by object name into six collections: **01_Character / 02_Prop / 03_Set / 04_Light / 05_Camera / 06_Particles**. A collection is both an organising unit and a render switch: with the same camera and the same lights, flipping collection visibility alone gives you a character pass, a prop pass and a set pass.

Element breakdown for the 11 3D shots (object counts per collection)

Shot	Frames	Character	Prop	Set	Light	Particles
S1 Highway	216	0	0	70	2	90
S2 Facing the city	192	29	2	115	8	90
S3 Inside the tower	144	0	0	35	3	70
S4 The interface	144	12	17	2	3	0
S5 Factory	96	20	11	42	7	70
S6 Corridor	144	19	2	29	6	46
S7 Whip pan	96	21	11	42	7	80

Shot	Frames	Character	Prop	Set	Light	Particles
S8 The climb	168	20	2	12	3	50
S9 The cell lands	144	0	3	27	3	0
S10 Taking it back	96	12	2	42	3	0
S11 Night city	144	0	0	117	32	0

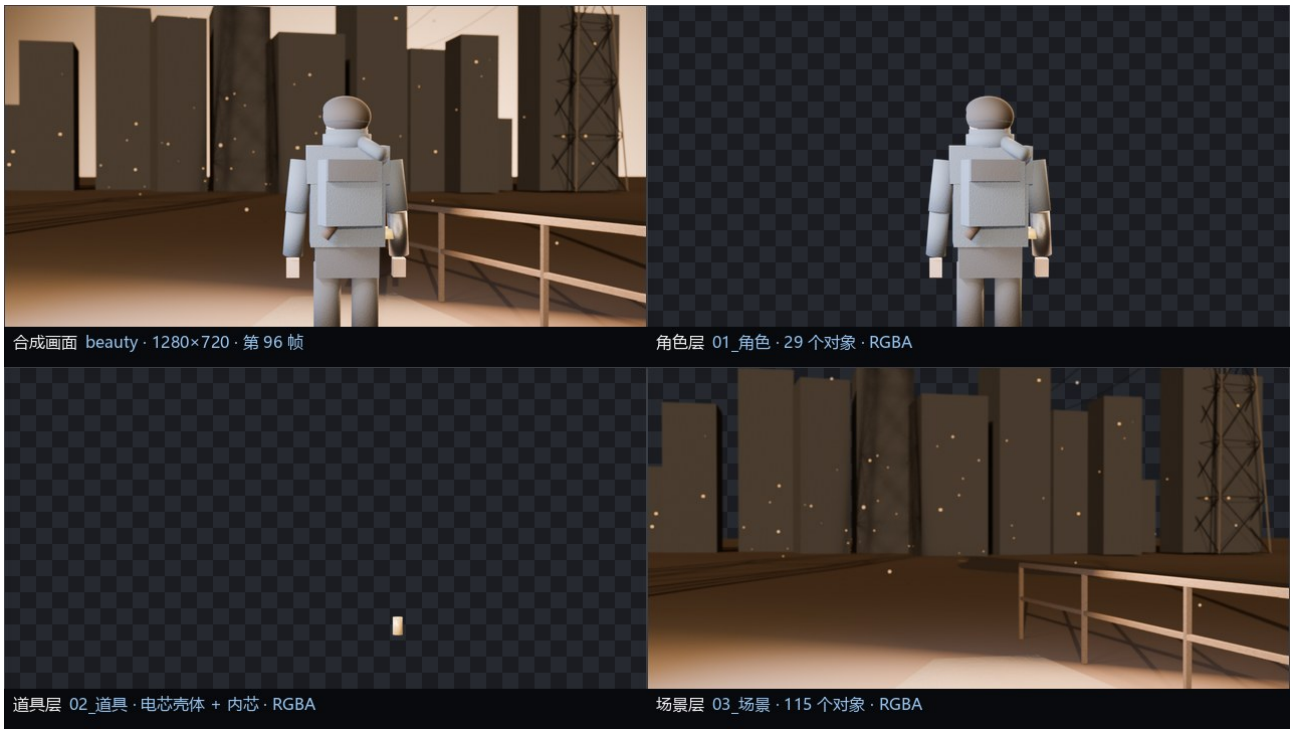
Totals: 1,584 frames, 133 character objects, 52 props, 533 set objects, 77 lights, 286,421 vertices / 552,610 triangles, and 11.5 MB across 11 project files. S1, S3 and S11 have zero character objects because those shots have no people; that is the point of the split, an empty layer stays empty instead of pretending. The 19 character objects in S6 are the patrol machine: it moves, turns and blocks the lead, so it belongs to the character collection, not the set.



The full matrix: 11 shots x 5 passes: composite, character, prop, set and clay. Each row comes from one frame of one shot with identical camera and lighting, so reading across shows one second split into layers. An empty cell is not a failed render: S1, S3 and S11 have no character, and none of those three has a prop either. The patrol machine appears in the character column of S6 and the cell appears in the prop column of S8; those two were the gaps this pass closed.



A heat map of object counts per collection, brighter meaning more objects. It doubles as a workload map: S2 and S11 carry the heaviest sets, S4 and S9 are almost all prop, and S8 has only 12 set objects yet has to carry the whole vault. Read it to predict which shot will break first.



Four passes of S2, all at frame 96: composite top left, character top right, prop bottom left, set bottom right. The character pass keeps the hand holding the cell intact, the prop pass holds only the cell shell and core, and the set pass keeps the distant city and road on its own. All four share one camera, so the composite needs no alignment.



The same shot in four panels: the composite top left, the character collection alone top right, the prop collection alone bottom left, the set collection alone bottom right. All four share one camera and one set of lights, so they stack back together seamlessly; the only difference is what is visible.



The S10 split demonstrated frame by frame: the picture cycles through composite, character, prop and set. Note that when the character is rendered alone the background is transparent - floor reflections and fog live in the set pass, not the character pass.

1. **Name first, then split into collections:** Object names are the classification rule, not your eyes. Give character parts a consistent prefix (head, torso, left arm, right arm, charge gauge), name props individually and name set pieces by structure (floor, walls, racks, pipes). If the names are a mess, fix the names before talking about passes; renaming costs far less than hand-drawing masks later.

2. **Group objects into six collections:** Character, prop, set, light, camera and particles each get their own collection with a name that says what it holds. Lights and the camera always stay visible - they are the shared truth about camera and lighting, and the moment each pass gets its own version, nothing will line up.

3. **Render a pass with visibility only:** Turn on a transparent background, leave only the target collection visible and render; then switch and render again. All passes must use the same frame range, resolution and colour management or they will drift out of register or shift colour. Transparent renders must carry alpha, otherwise the edges keep a rim of background colour.

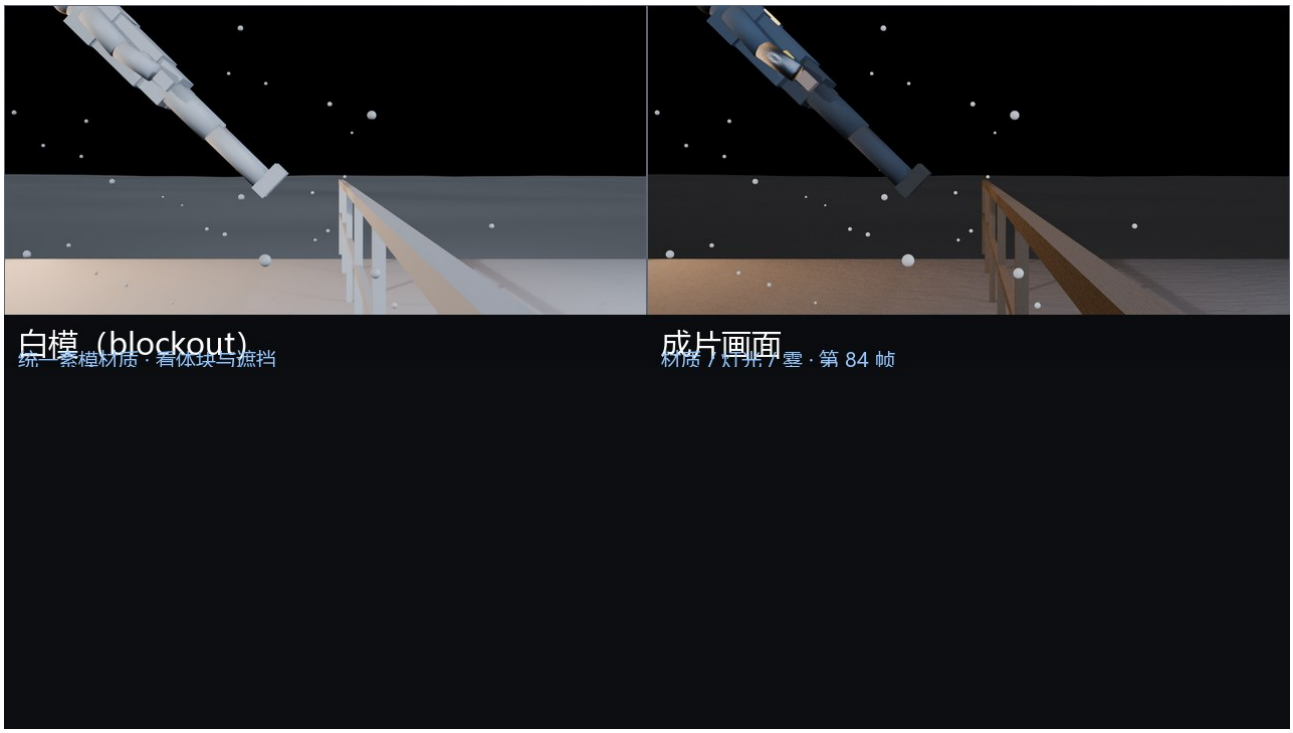
4. **Use object index for finer mattes:** To grade one part on its own, such as the mechanical arm or the charge gauge, give the object a stable object index and output a matte from that index; do not hand-draw roto on every frame. The index travels with the object, so the matte realigns after a re-render.

5. **Clay first, materials later:** Render the same shot once more in a uniform clay material: every surface becomes one neutral grey and only form and light remain. Clay instantly exposes scale errors, intersections and camera problems, far faster than hunting them on a shaded frame.

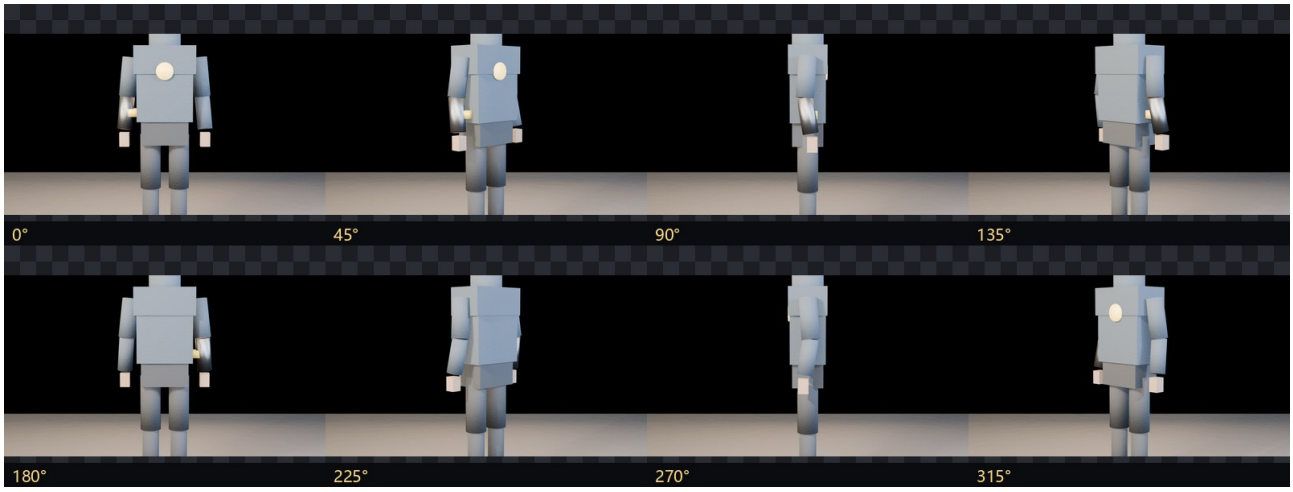
6. **Check characters and props with turntables:** Render a 360-degree turntable for the character and for each key prop: for the character check silhouette, centre of gravity and which arm carries the mechanism; for the prop check grip scale and material. A turntable turns only the subject with no set, which makes problems easiest to see.

7. **Grade the composite as one picture:** Bring the character and set passes back into the composite and unify exposure, fog, grain and vignette, so the figure does not look pasted onto a backdrop. Keep a version before and after the grade so any colour move can be undone.

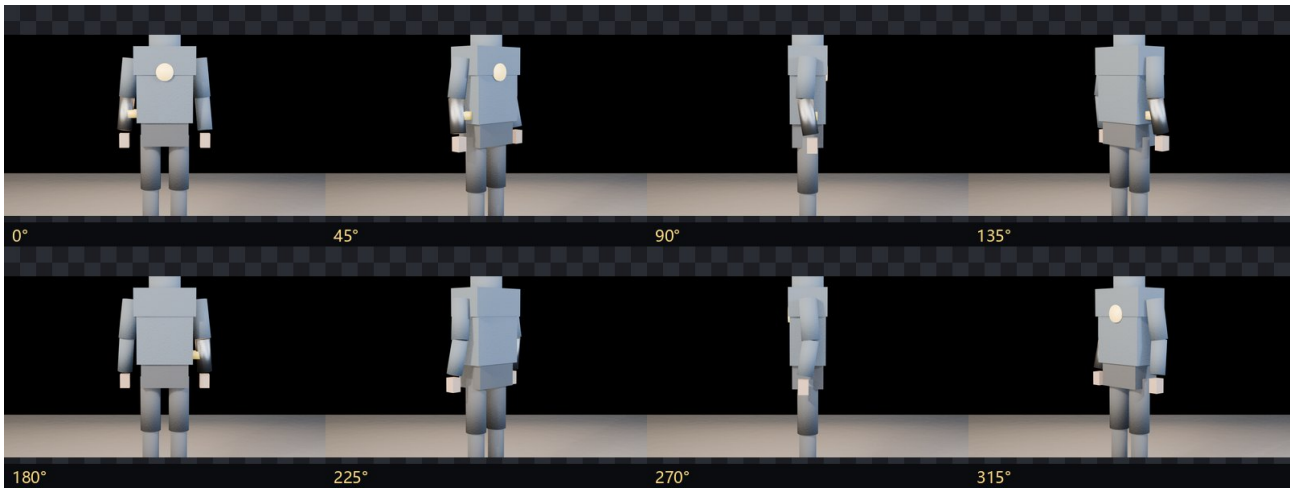
8. **Change one variable at a time:** Never change lighting and materials at the same time when re-rendering the character pass alone. Compare the same frame afterwards and confirm only the intended thing changed before moving to the next shot; otherwise you cannot tell which step introduced a problem.



Clay next to the finished frame. The left image keeps only form and light, the right image carries the final materials; camera, composition and pose must match exactly. A composition problem left unsolved in clay only gets more expensive once materials arrive.



Four frames of the character turntable: front, side, back and the other side. They verify that the mechanism sits on the left forearm, the charge gauge sits left of centre on the chest and the wear on the work clothes concentrates at knees and elbows - if any of those drift between shots, the audience drops out of the film.



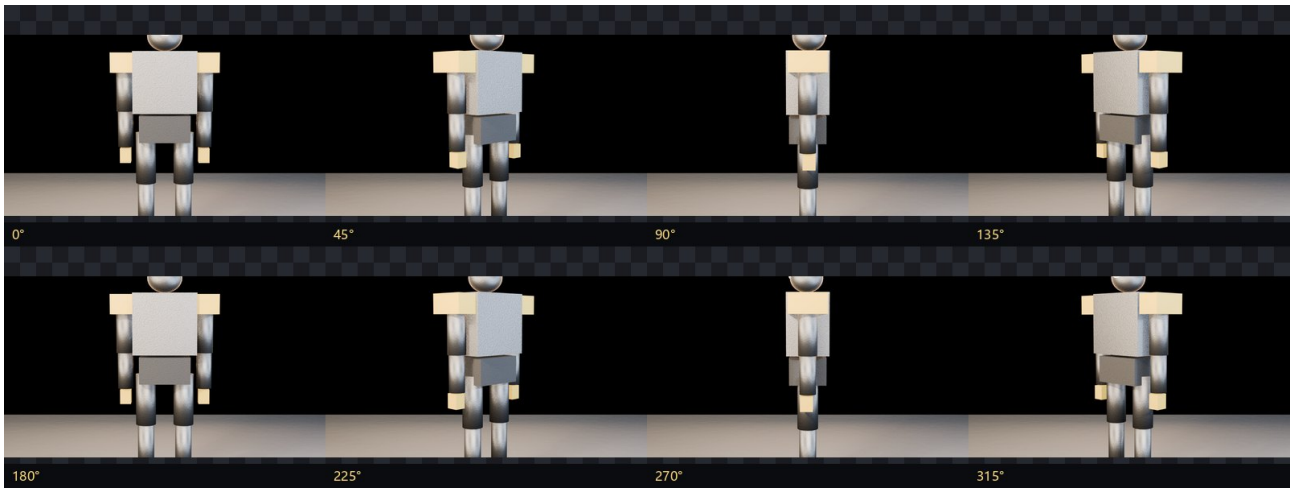
The full 96-frame character turntable (4 seconds at 24 fps). The turntable is the second safety net after the character card: the card is the intent, the turntable is what the model actually looks like.



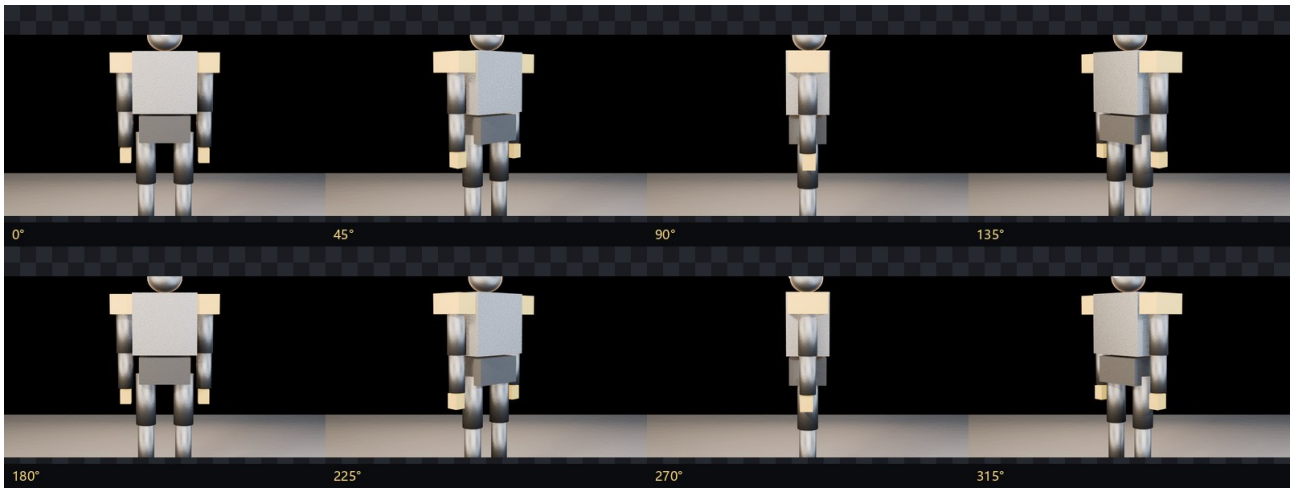
Prop turntable: the energy cell rotating once around its vertical axis. Watch the length-to-width ratio, wear at the grip point, whether the glowing core holds a steady brightness, and whether the shell breaks at any angle.



The full 96-frame prop turntable. Props drift between shots more easily than characters do, and a turntable is the cheapest way to pin their look down.



Eight frames of the patrol machine turnaround, one every 45 degrees. It has to read as a machine and still read as something that watches you in silhouette: head tipped forward, single eye set high, uneven shoulder line. Lock the silhouette on the turnaround before discussing materials; if the silhouette does not read as a character, better material only makes it a better device.



The full 96-frame turnaround of the patrol machine (4 seconds, 24 fps). It is on screen for only a few seconds in the S6 over-the-shoulder shot, but whether the audience feels watched depends on the model holding up from every angle.

Props must be parented to the character rig

Layering has a hidden trap: a prop living in its own pass does not mean it follows the hand. In S8 the lead hooks the cell in one hand and vaults the rail; an early version placed the cell in the prop collection and animated its world position by hand. As soon as the rig moved, hand and cell drifted apart. The character pass looked fine and the prop pass looked fine, but stacked together they broke.

The fix is to parent the cell object to the hand bone of the character rig so position, rotation and scale all inherit from the bone instead of being keyed in world space. The cell now follows wherever the character moves, and it stays in the 02_Props collection, so it can still be toggled and rendered on its own. Parenting solves following; collection membership solves separate

rendering; the two do not conflict.

1. **Find the hand bone first:** In edit or pose mode of the rig, confirm the name of the bone that drives the hand and write it down. Do not pick by eye; rename duplicate bones before you start.

2. **Parent to the bone, not a stack of constraints:** Select the cell, then the rig, enter pose mode, make the hand bone active and parent with keep transform. One parenting relation is enough; duplicating transforms and stacking three child-of constraints means nobody can say which layer is in charge when it breaks.

3. **Keep the collection membership:** After parenting, the cell still lives in 02_Props. Check the outliner: it should appear under the character rig and still appear in the prop collection, and both must be visible.

4. **Check three time points:** Render one composite frame at 0.2, 0.5 and 0.8, zoom in on where fingers meet the shell, then turn off the character pass and keep only the prop pass to confirm the cell path matches the hand path exactly.

After the fix: S7 and S8 were rebound the same way and their full sequences re-rendered; the prop objects in S8 went from 0 to 2 (shell and core) and the prop total in the element list went from 48 to 52. This kind of fault never shows in a single-layer still; it only appears when the two layers are stacked and the contact point is checked frame by frame.

Measured cost: The S10 passes are 96 frames each at 1280x720: character 59.7 s, prop 51.3 s, set 103.0 s, about 3.6 minutes in total; the 168-frame clay pass and both 96-frame turntables ran on the same machine. Splitting costs more renders and buys you the ability to fix any single layer without recomputing the whole shot - which is the practical answer to whether a character can be rendered on its own: yes, and it should be.

Layering is more than easier fixes. A separate character layer can be fed straight to a video model as the first frame, so the background is never re-imagined; the set layer can carry its own camera move for compositing or as reference for the next shot.

Step five: from previz frame to finished shot (optional refinement)

This master renders straight from the previz frames with no model repaint. If you want a few shots in a different texture, that is the optional branch below: once the previz frames are locked, generate shot by shot from the shot table, and use the first frame whenever you can. The first frame locks composition, character and lighting in one pass, so the model only has to add motion and detail; fall back to text-to-video only for shots with no usable previz frame. Any shot that does use refinement must be marked separately in the shot table.

> [Project] 75-second industrial sci-fi short, 16:9, photoreal cinematic, low-saturation teal and orange separation, light film grain, no text, no subtitles, no watermark, no logo.
> [Character] Shen Yan: adult man, lean, faded blue-grey workwear, bare metal on the left forearm with a band of warning yellow, amber charge gauge on the left chest; never shouts, never looks at the camera.
> [Set] Abandoned substation: high-window side backlight, wet floor reflections, rows of cell racks, yellow-black hazard lines; cool teal ambient, warm practicals.
> [Prop] Storage cell: 24 cm long, translucent amber shell, glowing core, one-handed; no deformation, not square, no cables.
> [This shot] S2 medium close 45 to 52 mm slow push: the figure stands with his back to camera outside the highway guardrail, looking at the dead-city skyline, a glowing cell in his hand; wind pushes the hem one way; side backlight 4300 K, cool rim.
> [Camera] Push 1.5 m along the line of sight, slow then quicker, stop at the end; no pan, no crane.
> [Consistency] Clothing, arm side, gauge position and cell appearance match the subject cards; one main action per shot; no extra people, text or watermark.
> [Negative] Face swap, arm on the wrong side, jacket colour shift, blown-out cell, camera shake, slow motion.

1. **Lock the look with a still first:** Share one look still across three consecutive shots; generate a first frame per shot, adopt it, then animate. Do not skip the first frame and go straight to text-to-video, or cross-shot consistency is pure luck.

2. **At least two candidates per shot:** Compare side by side and preview each one in full; adopt only when movement, figure and light all hold. A poster frame is not a preview.

3. **Write the adoption back into the shot list:** Record the adopted candidate ID in the table for rework and credits; keep rejected candidates on the canvas for comparison instead of deleting them.

4. **Quote before you confirm:** Generation is billed per task. Get a quote for the whole film first and submit only after checking the total; an expired quote must be re-quoted, never reused.

Step six: edit and sound

Editing does three things: order by the table, trim on action, and breathe at the two black gaps. Keep one black gap between S0 and S1 and another between S10 and S11, fading out six frames before and in four frames after; everything else is a hard cut landing three frames before the wind-up.

Sound has four layers: ambience (wind, ground rumble, metal expanding), action (footsteps, metal seating, landings), emotion (cell hum, isolated only in S9) and transitions (the breath of the two black gaps). There is no melodic score, only low pulses, which avoids licensing problems and keeps sentiment from drowning the action.

Do it yourself: rebuild this film in your workspace

1. **Create a blank canvas and state the boundaries:** Create a blank canvas and write the title, 75 seconds, 16:9, 13 shots and no dialogue into the brief. Build three subjects first: character Shen Yan, set abandoned substation, prop storage cell; upload one main image and at most two supporting images each.

2. **Write the five acts and the shot list:** Allocate time with the five-act table above, then fill shot size, camera, lens, movement, light, duration and source for each shot. One main action per shot; appearance lives on the subject card, action lives in the shot description.

3. **Do the 3D previz first:** Block the low-poly set, fix camera and movement, render character and set layers separately. Check first, middle and last frame per shot: no black frames, no jumps, brightness variation under 15 percent.

4. **Optional: drive the shot from its first frame:** For shots where you only want a different texture, generate video from the previz frame as the first frame, two candidates per shot. Preview in full before adopting, then write the candidate number back into the shot table. Shots you do not refine keep the rendered 3D result.

5. **Edit the film, quote and export:** Order the 13 shots, trim to the table, place the two black gaps; assemble the film manually in the platform, confirm the quote, export 480p, then re-encode a 720p master locally from the same timeline.

Completion criteria: The 13 shots play back to back with the character, the mechanical arm, the gauge and the battery continuous across cuts; the head and tail cards are legible; the 720p master and the 480p platform file both play from a direct link, with resolution, duration, audio and the table all agreeing; the page states the source of every shot: text card, 3D previz, and AI generation only where refinement was actually used. Uploaded assets alone, a single candidate per shot, or a queued export do not count as done.

Asset and cost boundaries

□ Inside the platform: The three subject cards, the 13-shot contact sheet, the layer comparison, the 75-second master and the platform 480p export are published with this tutorial for direct download and comparison.

□ Re-creation: Upload, canvas organisation, manual film, editing and 480p export are available; generation is billed per task, so read the quote before confirming and never submit an unauthorised paid task.

□ Remix: You can reshoot your own version from the shot list; if you reuse this project's characters and sets, keep the attribution so readers do not mistake them for your own designs.

To learn canvas basics first, go back to Case 2; to have an agent batch out shots, read Case 3; to train shot language systematically, start with Advanced creation.

Case 5: A complete nature-fantasy film • from script to The Spirit of Spring Mountain

What you see first: An 82-second, 12-shot nature-fantasy short, The Spirit of Spring Mountain, with links to a 720p master and the platform-exported 480p file, plus four subject turnarounds and one layered-render comparison clip. Every frame is built by this project from zero: the five-act causal chain and the 12-shot table first, then each character, prop and set, then a 3D previz, then character, prop and set rendered as separate passes, and finally one timeline, four sound layers and delivery. Compared with case 4 it adds natural sets, an animal character, a stone spirit, a ritual prop, four-season change and an environmental crowd of trees and flowers - good practice for making several subjects each hold up in the same space.

Studio → Canvas creation → New blank canvas → Five-act script → 12 shot nodes → Subject cards → 3D previz → Layered rendering → Manual film assembly → Export

Live project: This film is saved as a real project in the studio: the workflow project The Spirit of Spring Mountain holds the 82-second brief and the 12-shot table, and the canvas Spring Mountain · 12-shot production board holds 58 nodes - six process notes, twelve storyboard nodes, thirteen film segments, twelve previz stills, eight asset cards, four turnarounds and one layered comparison - feeding one film node. Sign in and open it for comparison: Spring Mountain · 12-shot production board The finished film was exported from the platform's manual edit at 854×480, 82.000 seconds, 24 fps, AAC stereo.

First the results: one local master, one platform export



If it will not play, download the 720p master (1280×720, 82.007 s, 24 fps, stereo).

The 82-second delivery of The Spirit of Spring Mountain: a 4-second title card, a 78-second main film (12 shots, 1872 frames) and a 4-second end card. Watch it twice: first for why she climbs the mountain, then for the cuts, noting at which second each space changes.



If it will not play, download the platform 480p export (854×480, 5.9 MB).

This copy is what the platform's manual edit → export produced from the same 13-segment timeline: 854×480, 82.000 seconds, AAC stereo, 5.9 MB, with no local re-encode. The 720p master was rendered locally from the identical edit, so cuts and audio match. Acceptance is judged on the delivered file's real resolution, duration and audio track.



The 12-shot thumbnail sheet in film order. Read the whole first: blue-grey cabin, jade-green forest path, cool-grey pass, blue-white summit, fresh green and pink-white return of spring. Each of the five spaces owns a primary colour, so viewers feel which act they are in without subtitles.



The 82-second timeline. S1, S2, S5 and S11 run 8 seconds each, the other main-film shots 6 seconds each, and both cards 4 seconds. S11 is the only uncut long take, running from 01:10 to 01:18; it carries the payoff, so it gets the longest continuous stretch.

Production chain: eight stages, each leaving a checkable artefact

From one sentence of story to a short film you can link to

Stage	What you do	Artifact and acceptance evidence
1 Script	Logline, theme, five-act causal chain, shootable scenes	The script section of the production bible; one sentence must say who lacks what, who blocks the way, and what she finally gains
2 Shot list	Shot size, camera, lens, movement, key light, duration, source for every shot	A 12-shot table; 1872 frames and 78 seconds in total, within 24 frames of the finished film
3 Subject cards	One card each for three characters, one key prop and four sets, fixing scale and material	Eight asset cards; the braid, antlers, spirit core and bell chip never drift between shots
4 Per-shot files	One file per shot, split into six collections: character, prop, set, lights, camera, particles	12 shot files and an element manifest; object, polygon and material counts are checkable per shot
5 Layered rendering	Character, prop, set and clay rendered as separate passes, sharing lights and camera	Pass sequences aligned frame for frame; switch off any layer and the rest still render on their own
6 Clay renders and turnarounds	Clay checks composition, turnarounds check silhouette and material	Clay comparison; four 96-frame turnarounds that hold up from every angle
7 Editing and sound	Order by the table, cut on action, lay four sound layers, two black gaps	An 82-second timeline; three bell rings span the cuts, each cut lands before the action starts

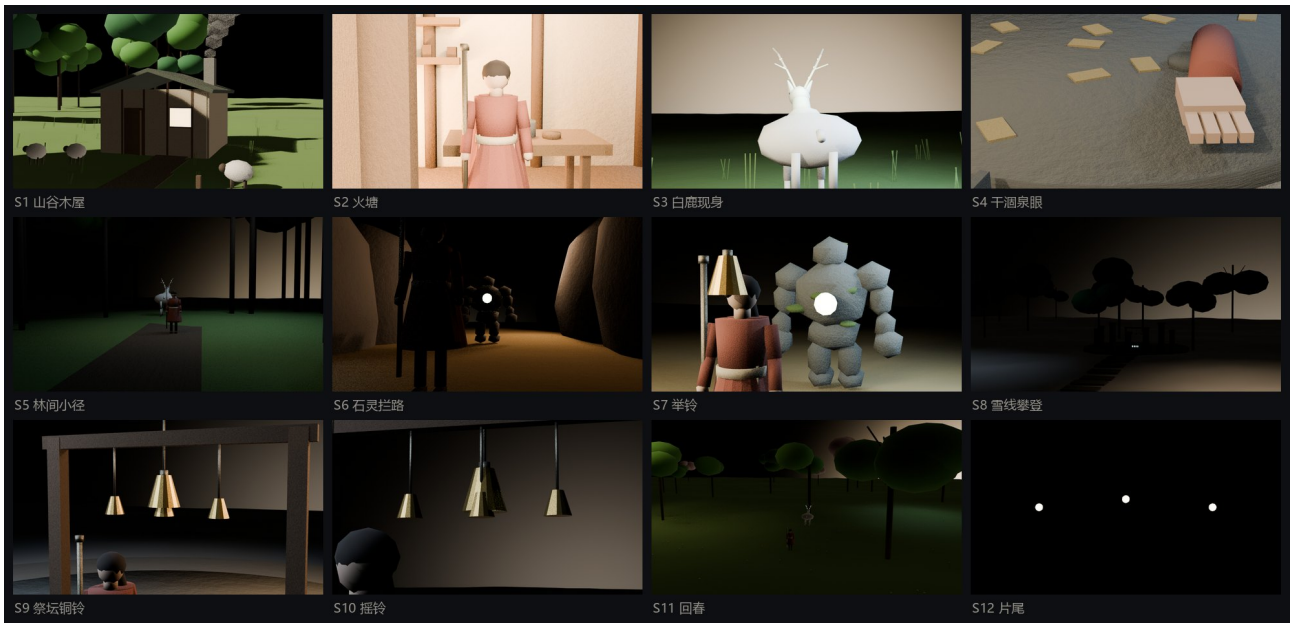
Stage	What you do	Artifact and acceptance evidence
8 Export	Platform-exported 480p plus local 720p master, from one timeline	Both files play from a direct link; resolution, duration and audio match the table

Step 1: write why she climbs the mountain before writing images

This film has no villain dialogue and no narration, so cause and effect must be written into action: **the spring runs dry → the white deer leads → the stone spirit blocks the way → the bell grants passage → the bell rings on the snow summit → spring water returns**. Each link maps to one visible action, and the audience needs no explanation - they just follow the road to the end. Fix these six actions first, then allocate time; write the timecodes as the final cut points so the structure never has to be reworked.

Five-act structure and causal chain (matches the 82-second film)

Act	Timecode	Function	Key action	emotion
0 Title	00:00-00:04	Set the tone, give the title	Black frame; title and subtitle fade up; a bell sounds far away	Still
1 Late spring	00:04-00:20	Establish that spring never came	Slow push on the valley; by the hearth, fingers trace the bell's chip	Cold, waiting
2 The deer leads	00:20-00:40	Give the goal and the route	The white deer appears; a palm presses the crack in the spring; she follows the deer into the forest	Curious, resolved
3 Stone barrier	00:40-00:52	Set the single obstacle	The stone spirit leans down; Ahe lowers her staff and raises the bell	Pressure, testing
4 Snow ritual	00:52-01:10	Complete the ritual	Climb the snow line; fingertips touch the bell rope; three rings	Isolation, solemnity
5 Spring returns	01:10-01:18	Pay off the promise	Water seeps from the same crack; blossoms open; the deer returns	Release, aftertaste
6 End card	01:18-01:22	Leave an aftertaste	Three warm-white points of light rise; the title fades	Still



The 12-shot 3D previz contact sheet. Previz answers only three things: are the spatial relations right, can the camera read the action, and does the light position separate the subject from the background. Materials and detail can come later; a wrong composition must be redone.

Step 2: the 12-shot table is the construction drawing

A shot table is not prose; it is a construction drawing you tick off item by item. Every shot lists at least seven things: framing, camera and lens, movement, light position and colour temperature, frame count, source, and the continuity anchor this shot must satisfy. The table below is the version this film actually used; its frames add up to 1872 and match the finished cut frame for frame.

The 12-shot construction table for The Spirit of Spring Mountain (78-second main film)

#	Length	Shot size	Camera / lens	Movement	Key light and colour temperature	Frames	Duration
S1	Valley cabin	Extreme wide to wide	Low angle from the slope, 38→47 mm	Push in 6 m, tilt up slightly	Backlit morning sun 4300 K, cool skylight fill	192	8 s
S2	Cabin hearth	Medium close to close-up	Eye level across the table, 38→46 mm	Push in 1.7 m	Warm hearth point light 2400 K, cool fill from the window	192	8 s
S3	The white deer appears	Medium	Behind and beside Ahe, 55→66 mm	Gentle push 1.2 m	Backlight in mist 5000 K, antlers self-lit in jade green	144	6 s

#	Length	Shot size	Camera / lens	Movement	Key light and colour temperature	Frames	Duration
S4	The dry spring	Overhead close-up	Above the spring, 65→80 mm	Press down and push in	Warm light from above the side 3600 K, low stone reflectance	144	6 s
S5	Forest path	Wide	Follow from behind, 32→40 mm	Follow 8 m	Canopy light 5200 K, ground mist	192	8 s
S6	The stone spirit bloc ks the way	Wide to medium	Low wide angle, 24→2 8 mm	Push in 1.3 m	Cool side light 6500 K, the spirit's cyan glow	144	6 s
S7	Raising the bell	Medium	Front side, 42→52 mm	Push in 1.0 m	Warm side light 3800 K, warm gold in the bell mouth	144	6 s
S8	Climbing the snow line	Extreme wide	High angle behind and beside, 52→64 mm	Follow and rise	High-altitude cool light 7500 K, warm rim	144	6 s
S9	The altar bell	Medium	Outside the altar, 44→56 mm	Orbit 25°, gentle push	Blue-white ambient, point light in the bell	144	6 s
S10	Ringing the bell	Medium close	Low angle, 48→64 mm	Pull back 1.5 m	Cool ambient, warm-gold shock ring	144	6 s
S11	Spring returns	Wide to extreme wide	High angle pull back, 28→36 mm	Pull back 9 m, rise 5 m	Low warm sun 3200 K, backlight	192	8 s
S12	End card	Title card	Front on, 50 mm	Static; point s of light brighten	Black with warm white	96	4 s

Continuity anchors: The bell's chip faces up in S2 and is still on the same side when she raises it in S7; the palm in S4 and the fingertips in S9 belong to the same hand, with the cuff worn in the same place; the white deer stays 4 to 6 m ahead of Ahe in S5; S6 and S7 share one space and one light position, changing only camera and framing; the water in S11 must seep from the crack in S4. Write these five beside the shot table - far more useful than the phrase "keep it consistent".

Step 3: turn every character, animal, spirit, prop and set into a card

This film has more subjects than case 4, so the cards must be harder. A card holds three things only: scale, material, and what must never drift. Adjectives cannot be checked - "gentle" is unverifiable, while "one braid over the left shoulder, tied with a cream cloth strip" can be checked.



Ahe: 16 to 18 years old, 1.62 m tall, russet coarse-cloth jacket and matching cloak, cream leg wraps, dark brown short boots; black hair in one braid over the left shoulder; a hemp cord at the waist carries the copper bell; a 1.45 m wooden staff in her right hand. Forbidden: the braid moving to the right shoulder, the cloak turning bright red, added earrings or weapons.



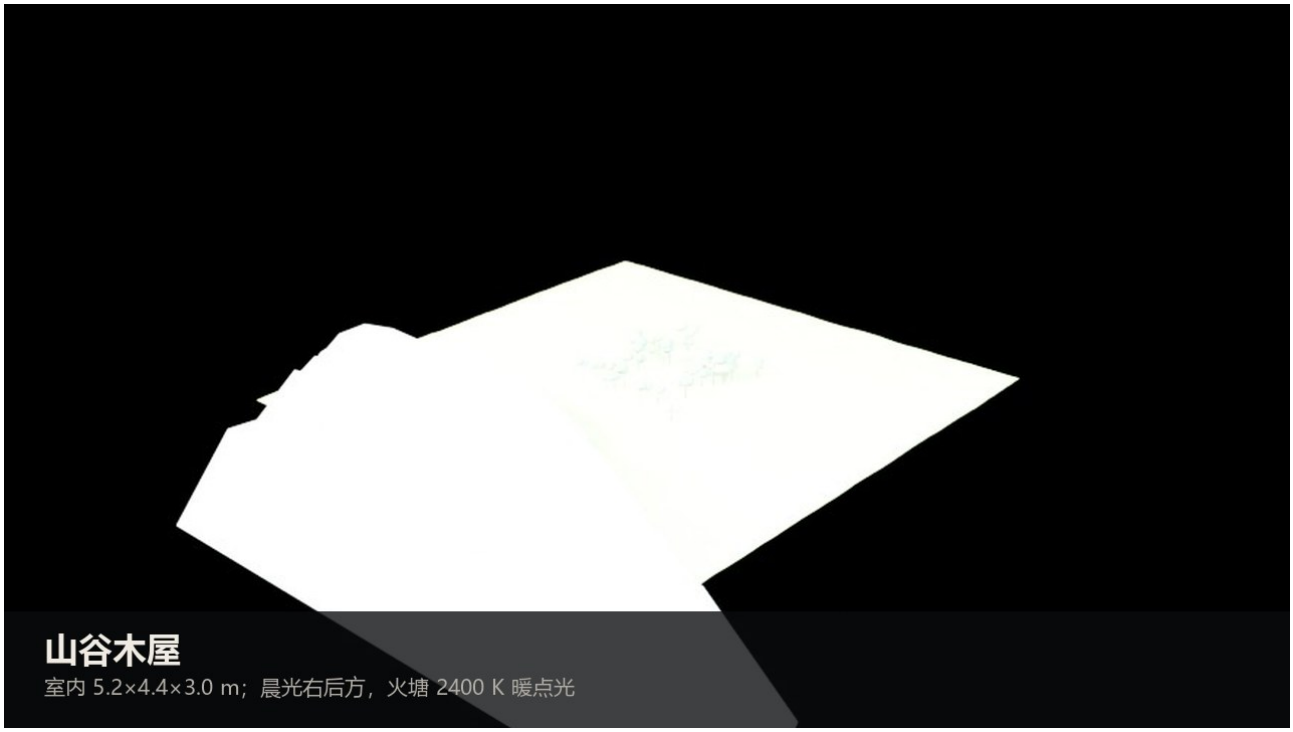
The white deer: shoulder height about 1.1 m, antler tips about 1.55 m, short white coat with grey around the muzzle and eyes, translucent antlers, five jade-green light lines along the flank. It is a guide, never cute or ingratiating; when it looks at Ahe it only turns its head and does not step closer. Forbidden: a nine-coloured deer, wings, or antlers branching into a tree.



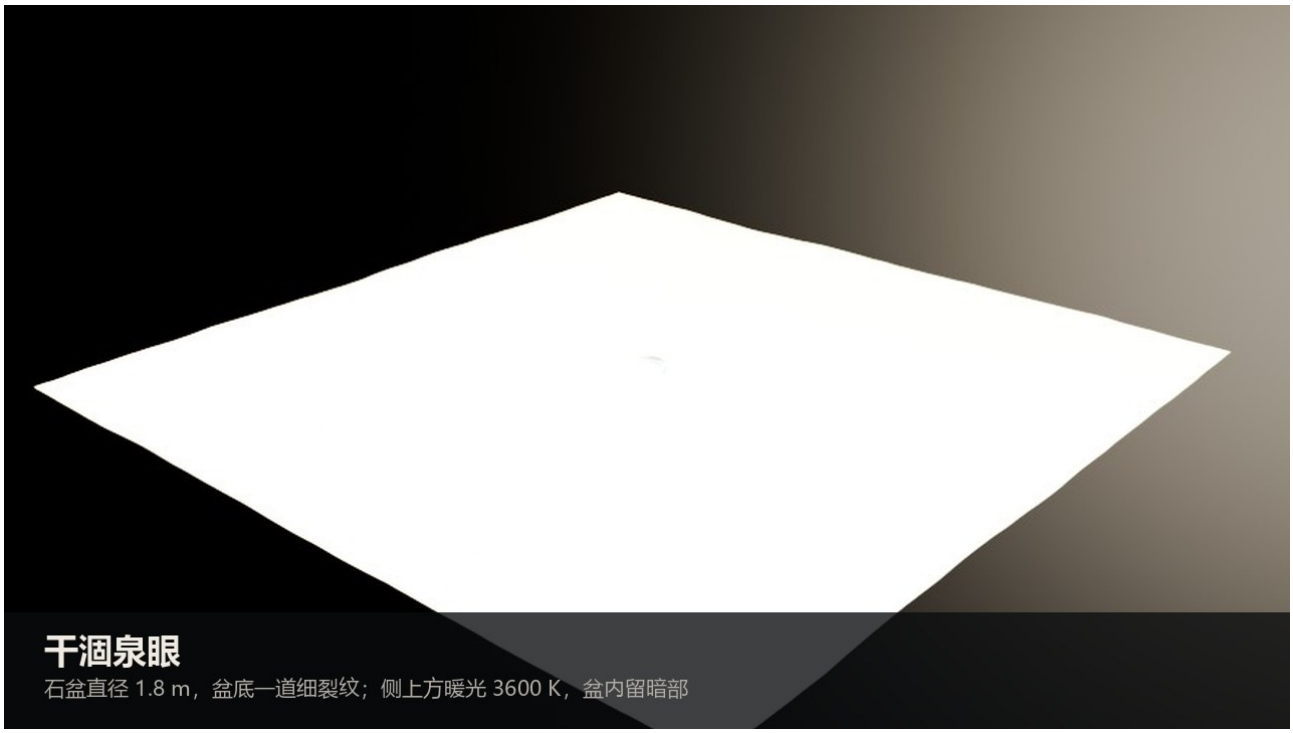
The stone spirit: about 2.3 m tall and 1.2 m across the shoulders, built from six to eight boulders, moss gathered on shoulders, back and knees, eyes and chest core in jade green. It rises in 3 seconds, leans down in 2, and steps aside only half a pace. It is not a villain but a gatekeeper: it never roars, punches or chases; the core never turns red and it has no metal joints.



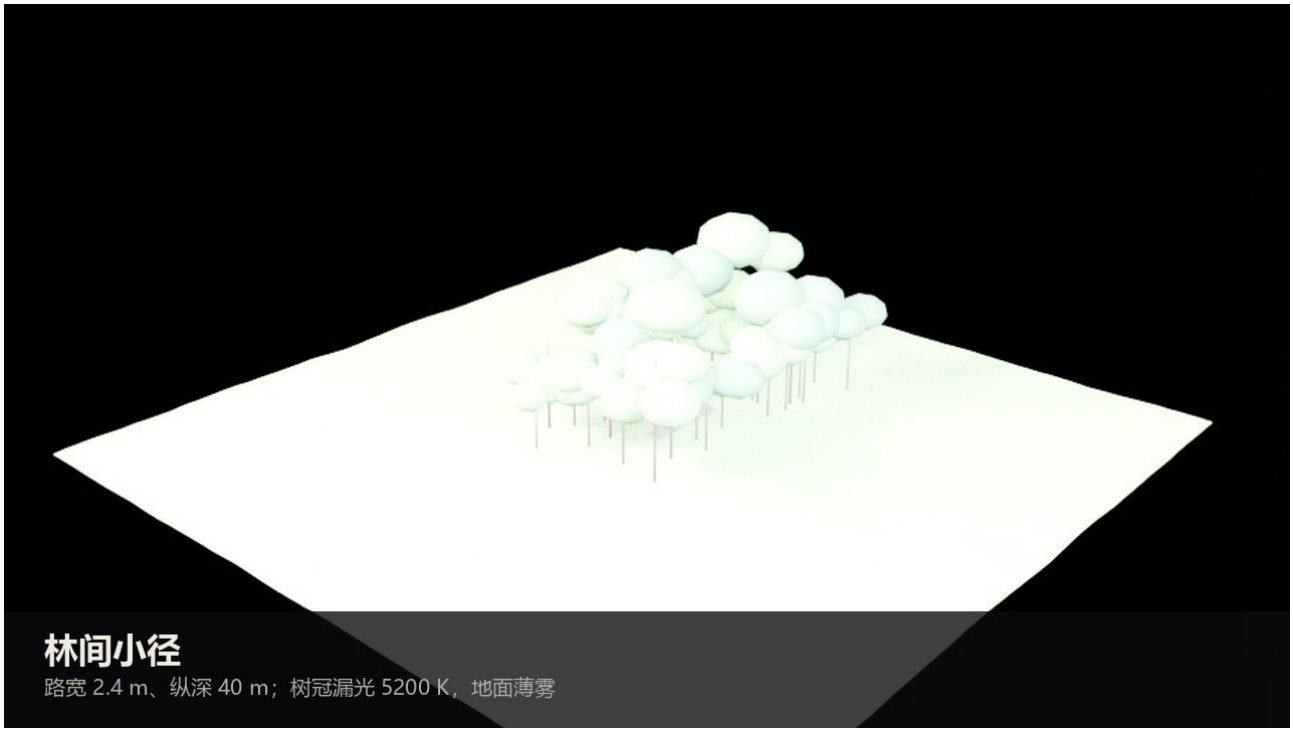
Copper bell: 19 cm high, 11 cm across the mouth, old bronze with green patina and small pits, a dark iron clapper, hung from hemp rope knotted at the crown. When it rings only the mouth glows warm gold and spreads in a ring; it never throws sparks, discharges electricity, or turns into crystal or a golden trophy.



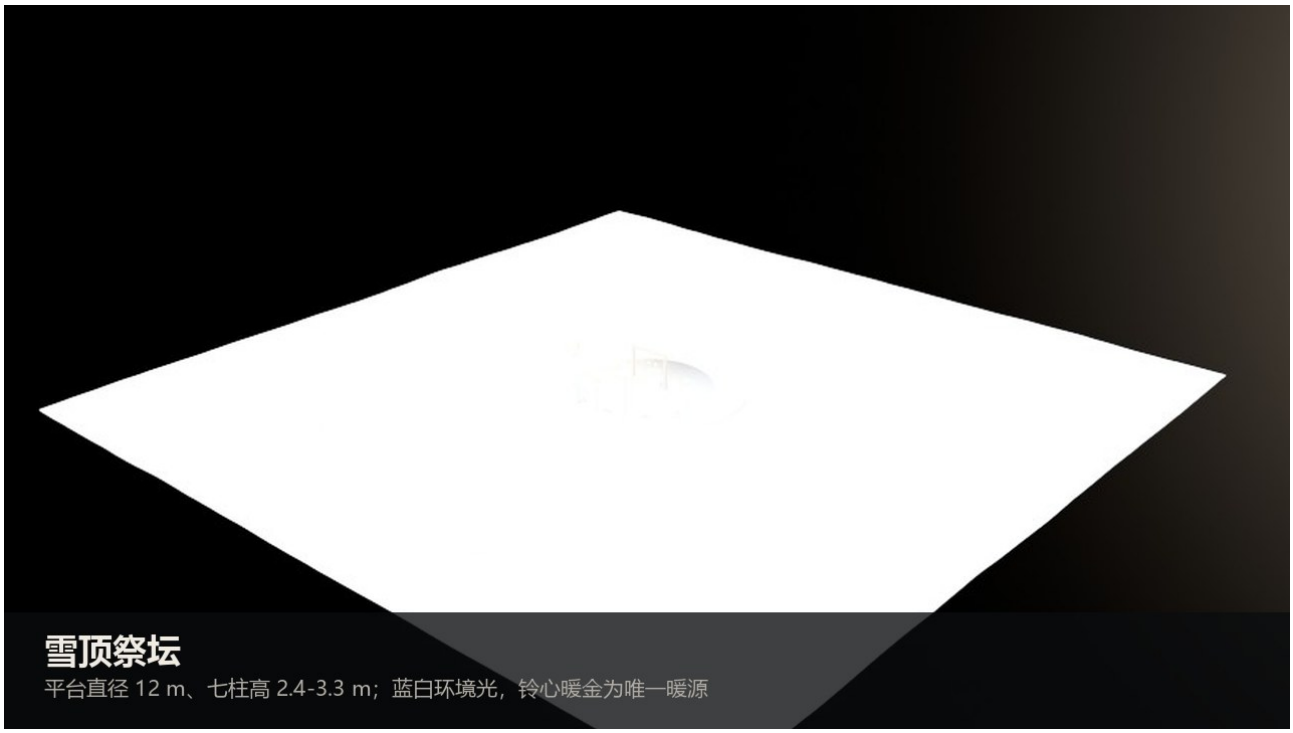
Valley cabin: 5.2×4.4×3.0 m inside, plaster walls, dark beams, old wooden floor, stone hearth; the door sits left of centre on the front wall, the window on the right, the chimney at the back right. Morning light comes from behind right; indoors the hearth is the only warm source.



The dry spring: a stone basin 1.8 m across, dead grass and rubble within 6 m; grey rock that has been wet and dried again, one fine crack across the bottom, dead leaves caught in the carved grooves. It is not an abandoned ruin but a place that has only just run dry - which is why the water in S11 must return here.



Forest path: 2.4 m wide and 40 m deep, trunks 9 to 16 m tall on both sides; wet earth, moss, fallen logs, ferns and thin ground mist. Dappled canopy light comes from upper left as key, with cool skylight filling from the right.



Snow-summit altar: a platform 12 m across, seven stone columns 2.4 to 3.3 m tall, a central stone platform 1.7 m wide, a bell frame 2.6 m high; snow, black rock, grey columns, moss and glowing carved grooves. Blue-white ambient dominates and the warm gold in the bell is the only warm source in frame.

Step 4: fix the colour script before talking about "cinematic"

The change of seasons does not come from twisting a colour filter; each act has its own primary colour, secondary colour and light position, so the audience changes breath with the cut. Tape this table beside the monitor and match colour before material when building a set.

Five-act colour script

Act	Space	Primary	Secondary	Light position	emotion
1 Late spring	Valley cabin	Desaturated blue-grey	Warm orange of the cabin	Backlit morning light, hearth indoors	Waiting, cold
2 The deer leads	Spring, forest path	Jade green	Amber light patches	Canopy light, backlight in mist	Curious, led
3 Stone barrier	Rock-wall pass	Cool grey	Moss green, the spirit's cyan glow	Low side light, hard shadows	Pressure, testing
4 Snow ritual	Snow line, altar	Blue-white	Warm gold of the bell	High-altitude cool light, point light in the bell	Isolation, solemnity
5 Spring returns	Valley	Fresh green, pink-white	Golden backlight	Low warm sun, aerial perspective	Release, aftertaste

Three hard rules: keep at least 10% of the snow summit in shadow and never blow snow to pure white; the highlight in the bell may clip locally but must stay under 1% of the frame; plants are never fluorescent green, always green with yellow-grey. Every exterior keeps aerial perspective, with distant ridges 1.5 to 2 stops lower in contrast than the foreground.

Step 5: one file per shot, split into six collections

One file per shot is where this pipeline actually saves time. Each file serves one shot but uses the same collection names internally: **01_Character / 02_Prop / 03_Set / 04_Light / 05_Camera / 06_Particles**. Consistent names let scripts toggle and render passes in batches; ad-hoc names start overwriting each other by shot three.



The 12-shot × six-collection element matrix. Read across to see how many objects of each kind a shot holds, and down to see whether one kind grows out of control across shots. S1 has the most set objects at 1039 because it is the establishing shot; S10 has exactly one particle object because the shock ring is its own effect layer.

Object and polygon counts for the 12 shots (from the element manifest)

Shot	Frames	Character	Prop	Set	Light	Particles	Vertices	Triangles	Materials
S1 Valley cabin	192	18	0	1039	2	0	53,183	101,326	16
S2 Hearth	192	25	4	44	4	0	1,350	2,410	15
S3 The white deer appears	144	43	0	271	3	0	13,311	24,972	5

Shot	Frames	Character	Prop	Set	Light	Particles	Vertices	Triangles	Materials
S4 The dry spring	144	6	12	17	2	0	8,729	16,960	5
S5 Forest path	192	67	0	212	2	0	24,019	46,408	15
S6 The stone spirit blocks the way	144	55	0	11	3	0	13,919	27,144	11
S7 Raising the bell	144	55	3	1	3	0	11,447	22,268	11
S8 Climbing the snow line	144	25	0	184	2	0	27,365	53,300	14
S9 The altar bell	144	25	9	59	3	0	12,089	23,412	12
S10 Ringing the bell	144	25	3	59	2	1	8,889	17,116	13
S11 Spring returns	192	68	0	557	2	0	42,999	82,778	15
S12 End card	96	0	0	4	1	0	134	252	2

Totals: 1872 frames, 217,434 vertices, 418,346 triangles. The point of these counts is to catch two problems early: a shot with ten times the polygons of its neighbours has not cleaned up its set, and a shot with far fewer materials than expected has merged things that should stay distinct. S4 once counted fingers as props, giving character 0 and prop 18; corrected, it reads character 6, prop 12, set 17, with the fingers moving with the palm and only the basin and carved grooves left in props.

Step 6: split characters, sets and props into separate renders

Yes, and you should. Layering is not slicing the frame into three pictures; it requires three passes to line up exactly **at the same camera, the same lights and the same frame number**. The method: keep every object in one scene and change only collection membership; render each layer with view layers or collection toggles. Lights and cameras stay visible in every layer, otherwise each pass lights itself differently and no amount of compositing will line them up.



Four panels: character pass top left, prop pass top right, set pass bottom left, final composite bottom right. Note where the hand meets the bell in the character pass; it must coincide exactly with the composite. The prop pass keeps only what a character can hold or push - never the whole forest.



Layered comparison clip: six time points, four panels switching in sync. Watch three things - whether the character pass drags the background with it, whether a prop leaves the hand as the character moves, and whether the set pass drifts with the camera.

1. **Define collections, not parenting:** Put each shot's characters, props and sets into collections 01/02/03. Collections answer "can this render on its own"; parenting answers "does this move with that". Never mix the two.

2. **Bind held props to bones:** Props carried by the hand, such as the bell and the staff, take the matching hand bone as parent (keep transform) and still stay in 02_props. Keyframe world coordinates alone and the prop leaves the hand the moment the character moves.

3. **Never split lights or camera:** 04_lights and 05_camera stay visible in every layer. They are not layers but the shared coordinate system; split the lights too and each pass is lit from scratch.

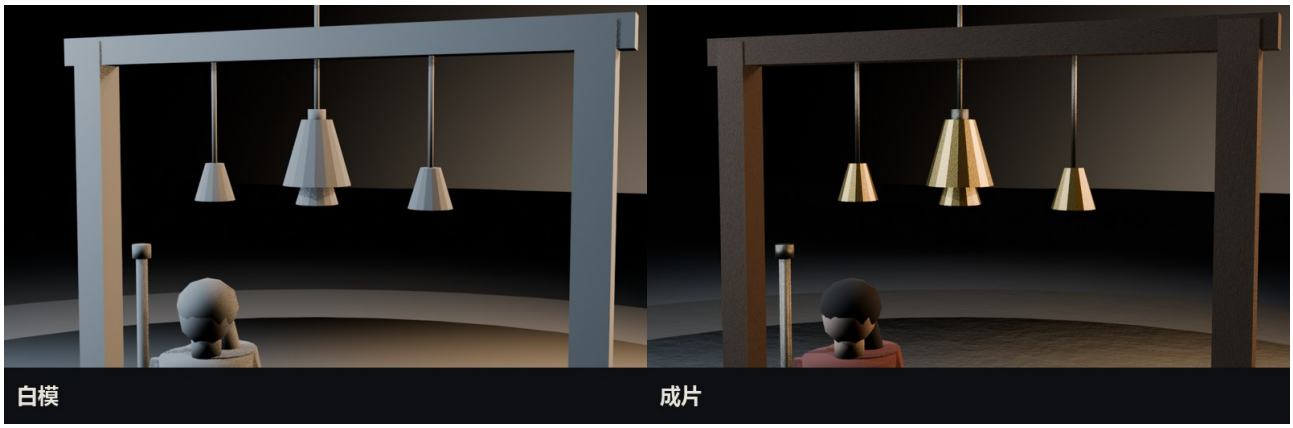
4. **Check on the same frame:** Pick any frame, export character, prop and set as three images with alpha, stack them in a compositor and compare with the final frame pixel by pixel. Edges, contact points and shadows must all line up.

5. **Add layers only when needed:** This film splits only character, prop, set and clay. Particles become their own layer only in S10, because the shock ring needs independent brightness control; every extra layer makes alignment harder.

> **Measured:** All 12 shots rendered at 1280×720, 24 fps and 20 samples - 1872 frames in about 9 minutes 40 seconds; S6, S9 and S10 also output character, prop and set passes, six verification frames each for alignment. Layering costs more renders and buys independence: change an expression and only the character pass is re-rendered; change the fog and only the set pass.

Step 7: clay renders check composition, turnarounds check silhouette

A clay render keeps only form and light, answering "can this camera read the action". A turnaround answers "does this subject hold up from every angle". Both come before final rendering, where they cost least.



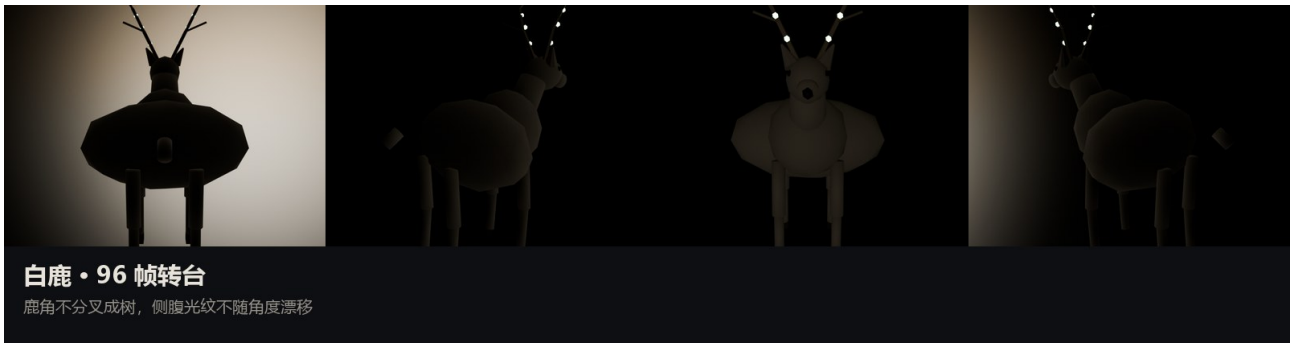
Clay and final side by side. The left image keeps only form and light, the right is the final material; camera, composition and pose must be identical. Any composition problem left at the clay stage only gets more expensive once materials are in.



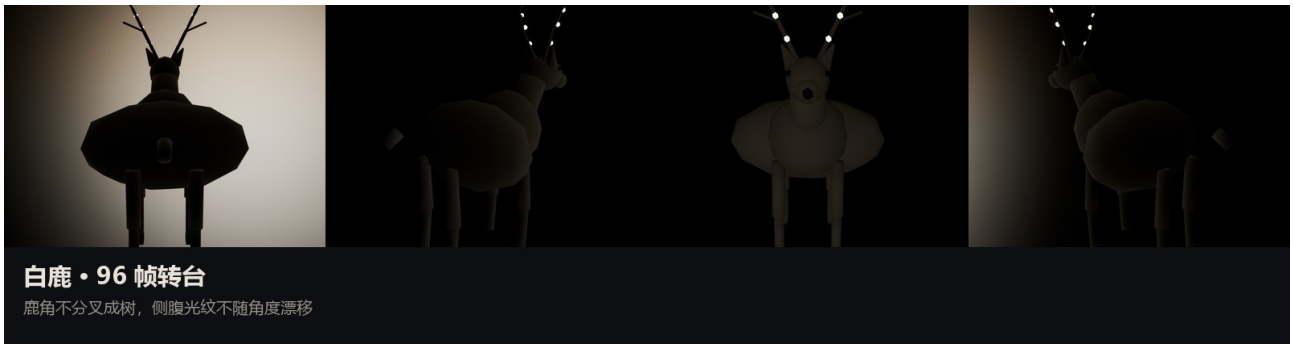
Four frames of Ahe's turnaround: front, side, back, other side. Check three things - the braid stays on the left shoulder, the cloak hem keeps its length, and the bell at the waist does not cut into the cloth.



Ahe's full 96-frame turnaround (4 s, 24 fps). The card is the intent, the turnaround is what the model actually looks like; when they disagree, fix the card.



Four frames of the deer's turnaround. Watch the antler branching, whether the five jade-green flank lines stay continuous through all four angles, and whether the legs look too thin in profile. Animal characters reveal proportion problems the moment they turn.



The deer's full 96-frame turnaround. It appears in S3, S5, S8 and S11; if any one angle fails, viewers will feel the deer was swapped in some shot.



Four frames of the spirit's turnaround. Read the silhouette: is the shoulder line wide enough, do the arms hang near the ground, does the chest core stay visible in profile. It appears in only two shots but must read "you cannot pass" at a glance.



The spirit's full 96-frame turnaround. Slow, heavy, sparing is its performance rule, and the turnaround must be slow too; spinning it fast makes it read as a machine rather than stone.



Four frames of the bell's turnaround. Check the position of the chip in the mouth, the patina distribution, and the gap between clapper and wall. The bell is the film's one key prop and must match in S2, S7, S9 and S10.



The bell's full 96-frame turnaround. Props drift more easily than characters from shot to shot; a turnaround is the cheapest way to nail down appearance.

Step 8: a reusable asset library

Eight asset packs are saved independently and shot files only link to them, never copy. The deer from S3 enters S5 unchanged and the bell from S9 enters S10 unchanged, so material and scale are never edited twice. Each pack keeps its original object and material names plus a preview; characters and props also get 96-frame turnarounds.

Asset packs and where they are reused

Asset ID	Type	Source shot	Contents	Reused in
char_ahe	Character	S2	Full body, braid, cloak, staff, bell at the waist	S1-S11
deer_spirit	Character	S3	Full body, antlers, flank light lines	S3、S5、S8、S11
golem_stone	Character	S6	Boulder body, moss, eyes and core	S6、S7
prop_bell	Prop	S9	Bell body, clapper, hemp rope	S2、S7、S9、S10
set_valley_house	Set	S1	Cabin, hearth, furniture, chimney	S1、S2
set_spring_eye	Set	S4	Stone basin, crack, carved grooves, dead leaves	S4、S11
set_forest_pass	Set	S5	Tall trees, earth track, mist, fallen logs	S5
set_snow_shrine	Set	S9	Altar, seven columns, stone steps, snow	S8、S9、S10

Reuse rule: When reusing a pack, change only the camera, lights and character pose, never the asset itself; if it truly must change, change it once in the pack and re-link every shot. If one shot needs a variant (say the deer must run in S11), copy it as `deer_spirit_run` and note the source in the manifest instead of editing the original pack.

Step 9: optional refinement - give a few shots a different texture

The master does not depend on model repainting; all 12 shots come from this project's 3D renders. If you do want another texture, take this optional branch: use the previz frame as the first frame to lock composition, character and light position, and let the model add motion and detail only. Any refined shot must be marked in the shot table; unrefined shots keep their rendered original.

[Project] 82-second nature-fantasy short, 16:9, realistic cinematic texture, desaturated natural palette, light film grain, no text, no subtitles, no watermark, no logo.
[Character] Ahe: girl of 16 to 18, 1.62 m tall, russet coarse-cloth jacket and matching cloak, cream leg wraps, dark brown short boots, black hair in one braid over the left shoulder tied with a cream cloth strip; a 1.45 m wooden staff in her right hand and an old bronze bell on a hemp cord at her waist. She never roars, glares or looks at the camera.

[Character] White deer: shoulder height 1.1 m, short white coat, grey around muzzle and eyes, translucent antlers, five jade-green light lines along the flank; it only turns its head and never steps closer or acts cute.
[Set] Snow-summit altar: platform 12 m across, seven grey stone columns 2.4 to 3.3 m tall, central stone platform 1.7 m wide, bell frame 2.6 m high; snow keeps its shadows, blue-white ambient light, warm gold in the bell is the only warm source.
[Prop] Copper bell: 19 cm high, 11 cm across the mouth, old bronze with green patina, dark iron clapper, hung on hemp rope; when it rings only the mouth glows warm gold and spreads in a ring, no sparks, no electricity.
[This shot] S10 medium close-up 48 to 64 mm: Ahe grips the bell rope with both hands, pulls back and releases; the bell rings for the first time and a warm-gold ring of light leaves the mouth and spreads across the snow; she neither shouts nor smiles, she only watches the bell.
[Movement] Camera at low angle pulls back 1.5 m at even speed and stops; no pan, no crane, no orbit.
[Consistency] Braid on the left shoulder, staff in the right hand, bell chip on the same side; snow-line height continuous with the altar floor; one main action per shot; no extra people, text or watermark.
[Negative] Face swap, braid switching shoulders, cloak turning bright red, antlers branching into a tree, metal joints on the spirit, bell turning to crystal, sparks or electricity, camera shake, slow motion.

1. **Lock the look with a first frame:** Consecutive shots of the same character share one look image; generate a first frame per shot, adopt it, then generate video. Never skip the first frame and go straight to text-to-video.

2. **At least two candidates per shot:** Compare side by side and preview each one in full; adopt only when movement, figure and light all hold. A poster frame is not a preview.

3. **Write the adoption back into the shot list:** Record the adopted candidate number in the table and mark the shot as refined; keep unused candidates on the canvas for comparison instead of deleting them.

4. **Quote before you confirm:** Generation is billed per task. Get a quote for the whole film first and submit only after checking the total; an expired quote must be re-quoted, never reused.

Step 10: cut for three things, mix four sound layers

The edit does three things: order by the table, trim on action, and breathe at the black gaps. The film has only two black gaps - title into the main film, and before the end card; everything else is a hard cut landing 3 frames before the action starts. S11 is not cut, letting the audience come down from the summit with the camera.

The 82-second delivery timeline

Segment	In	Out	Content	Duration
C0	00:00	00:04	Title card	4 s
S1	00:04	00:12	Valley cabin	8 s
S2	00:12	00:20	Hearth and bell	8 s
S3	00:20	00:26	The white deer appears	6 s
S4	00:26	00:32	The dry spring	6 s
S5	00:32	00:40	Forest path	8 s
S6	00:40	00:46	The stone spirit blocks the way	6 s
S7	00:46	00:52	Bell grants passage	6 s
S8	00:52	00:58	Climbing the snow line	6 s
S9	00:58	01:04	The altar bell	6 s
S10	01:04	01:10	Three rings	6 s
S11	01:10	01:18	Long take: spring returns	8 s
C12	01:18	01:22	End card	4 s

Four sound layers: ambience (valley wind, eaves, forest birds, summit wind), actions (footsteps, staff taps, cloth, stone, snow grains), emotion (hearth crackle, deer hooves, metal vibration) and transitions (bell tails, the half-beat of still air, low-frequency push). The three bell rings are the only loud events: the first is close, the second adds valley echo, the third keeps only low frequency and tail. There is no melodic score; space is worth more than sentiment.

Do it yourself: rebuild this film in your workspace

1. **Create the blank canvas and fix its boundaries:** Create a blank canvas and write the title, 82 seconds, 16:9, 12 shots and no dialogue into the brief. Build eight subjects first - Ahe, the white deer, the stone spirit, the copper bell, the cabin, the spring, the forest path and the altar - each with one main image and up to two supporting images.

2. **Write the five acts and the 12-shot table:** Allocate time from the five-act and shot tables above, then fill framing, camera, lens, movement, light position and duration for every shot. One main action per shot; appearance lives on the card, action in the shot description.

3. **Build 3D previz:** Block the low-poly set, then fix camera, lens and movement before light positions. Render first, middle and last frames per shot, checking for black frames or jumps and keeping brightness variation within one shot under 15 percent.

4. **Split the file into six collections:** One file per shot, with objects in 01_characters / 02_props / 03_sets / 04_lights / 05_camera / 06_effects. Parent held props to bones but keep them in the prop collection.

5. **Render four passes:** Render character, prop, set and a white-model pass separately, keeping light and camera visible in every layer. Composite three layers back on one frame and check edges, contact points and cast shadows.

6. **Pass clay and turnarounds first:** Clay confirms composition and four turnarounds confirm silhouette and material. A subject that fails its turnaround does not go into final rendering.

7. **Edit and lay the sound:** Order the 13 sections on the timeline, use black gaps only at the head and tail, and hard cut the rest. Lay ambience first, then action sound, then the three bell rings; never start with music.

8. **Export and check every field:** Export 480p through Manual film, then produce the 720p master from the same sequence. Check resolution, duration, frame count, audio track and file size, then actually play both files.

Completion criteria: The film plays through all 12 shots; Ahe's braid, staff and bell stay consistent across cuts; the deer's antlers and flank lines never clip; the stone spirit's rock walls, light and ground stay continuous between S6 and S7; and the water in S11 seeps from the same crack seen in S4. The character, prop and set passes composite back into one frame with matching edges and contact points, and all four turnarounds complete a full circle. The platform-exported 480p and the local 720p master both stream directly, at 82.007 seconds, 24 fps, with a stereo track. Uploaded assets, a single-shot candidate or a queued export do not count as done.

Asset and cost boundaries

□ Inside the platform: Eight asset cards, the 12-shot previz sheet, the six-collection element matrix, four turnarounds, the layered comparison, the 82-second master and the platform-exported 480p are all published with this tutorial for direct download.

□ Re-creation: Uploading, canvas organisation, storyboard structure, manual film, editing and export are available; generation models are billed per job, so check the quote before confirming and do not submit paid jobs that were not authorised.

□ Remix: You can reshoot your own version from the shot list; if you reuse this project's characters and sets, keep the attribution so readers do not mistake them for your own designs.

To learn canvas basics first, go back to Case 2; to have an agent batch out shots, read Case 3; to train shot language systematically, start with Advanced creation.

Case 6: A complete snow-epic production - from script to The Bell on the Snow Ridge

What you see first: An 87-second, 12-shot snow-epic short, "Ridge Bell", with a 1280x720 master, a platform-exported 480p cut, a 12-shot contact sheet and an 87-second timeline. Warm colour enters from only two places - an old bronze camel bell with a chipped mouth (its heart lights from S10 on) and the warm orange behind the cloud gap in S12; everything else sits in cold blue-grey. This film is colder than cases 4 and 5: no dialogue, no score, no villain lines, just a white camel, a snow fox, an Andean condor and one bell carrying loss, recovery, obstruction, the climb and the return. It drills the full 'few characters, few lines, hard light' workflow: script, shot list, subject cards, six-collection breakdown, layered rendering, AI retouch, four-layer sound and platform assembly.

Studio -> Canvas -> New blank canvas -> Five-act causal chain -> 12 shot nodes -> Subject cards -> 3D previz -> Layered independent rendering -> Manual film -> Export

Live project: The film is saved as a formal project in the studio: the workflow project The Bell on the Snow Ridge holds the 87-second brief and the 12-shot list, and the canvas The Bell on the Snow Ridge - 12-shot complete pipeline ties six process cards, 12 storyboard nodes, seven subject cards, five turnaround sets, layer comparisons, 12 shot clips, 13 refined stills and the film node into one readable pipeline. Open it after signing in: The Bell on the Snow Ridge - 12-shot complete pipeline canvas. The finished film was hand-assembled in the platform editor and exported from there; the timeline is one video track with 14 clips plus one audio track.

See the result first: one master, one platform export



If playback fails, download the 720p master (1280x720, 87.000 s, 24 fps, AAC stereo, 39.2 MB).

The 87-second delivery of The Bell on the Snow Ridge. The structure is a 4-second title card, 79 seconds of main film (12 shots, 1896 frames) and a 4-second end card. Watch it twice: the first pass only asks where the bell is, the second counts the three rings and notes which second each lands on.



If playback fails, download the platform 480p export (854x480, 5.0 MB).

This one is the same 87-second timeline pushed through Manual film -> Export in the platform: 854x480, 87.000 s, 24 fps, AAC stereo, 5.0 MB. Cuts and audio match the master, and both audio tracks line up sample for sample with the 87-second mix; acceptance checks the real resolution, duration and audio track rather than appearance or filename.



Twelve stills in film order. Look first at the warm-cool spread: the gate, the slope and the valley are almost entirely cold blue-grey, and warm colour is added a step at a time - the bell body only catches a reflection in S4, the bell heart lights warm gold in S10, and the warm orange only floods the snow in S12 (roughly 1.6% / 2.8% / 8.4% of pixels with R-B above 25). The viewer can feel when the light starts to warm up without any subtitle.



The 87-second timeline. Title and end cards take 4 seconds each; among the 12 shots, S1 and S8 run 8 seconds, S9, S11 and S12 run 7, and the rest run 6. The 7 seconds of S11 carry the payoff - the ring travelling over the valley - and it is the only long take in which the halo crosses the whole frame.

One-sentence story and the five-act causal chain

One sentence: **On a snowy night the pack bell falls from the gate lintel; a white llama and a snow fox chase it up the slope and hang it back on the lintel before dawn, and the first ring lights the whole valley.** With no dialogue, the causality has to live in action: **rope snaps -> bell falls -> chase -> condor bears down -> three animals together -> climb the spine -> hang the bell -> the ring returns to the mountains**. Each link is a visible action, so the audience never needs an explanation.

Five-act structure and causal chain (matches the 87-second film)

Act	Timecode	Function	Key action	emotion
0 Title	00:00-00:04	Set the tone, give the title	Black frame; title and subtitle fade up; a bell sounds far away	Still
1 Lost bell	00:04-00:18	Establish that the bell is gone	The white llama and the snow fox outside the gate; the rope wears through on the stone edge and the bell rolls down the slope	Cold, off balance
2 Pursuit	00:18-00:36	Give the goal and its cost	The fox looks back, drags the bell by the rope, pushes into the wind	Urgent, stubborn
3 Obstructio n	00:36-00:48	Set the single obstacle	The condor spreads its wings and bears down; the three animals share one shelter for the first time	Pressure, pause
4 Summit	00:48-01:03	Spend all the strength	The three stand apart on the spine; the llama bites the rope and hauls the bell up inch by inch	Isolated, gritted
5 Return	01:03-01:23	Pay off the promise	The three hang the bell together; three rings carry the halo over the gate and the valley; dawn closes the film	Release, aftertaste
6 End card	01:23-01:27	Leave an aftertaste	Snow turns gold, title fades	Still

Why only one bell: A short film fails when every shot looks good but the audience cannot tell what they are watching. This bell carries three jobs at once: it belongs to the llama (motive), it is the only warm light source (visual anchor), and it produces the three rings (audio anchor). Compress those three into one prop and 12 shots will not scatter. When you write your own film, find that one prop with three jobs before you storyboard.

The 12-shot list is the construction drawing

A shot list holds no adjectives, only parameters you can tick off: shot size, camera and lens, movement, key light and colour temperature, frames, duration. The table below is the version this film actually used - 1896 frames and 79 seconds of main film, frame-accurate against the finished cut; with 96 frames each for title and end cards, the total is 2088 frames and 87 seconds.

Shot list for The Bell on the Snow Ridge (79 seconds of main film)

#	Length	Shot size	Camera / lens	Movement	Key light and colour temperature	Frames	Duration
S1	Outside the snow-ridge gate	Extreme wide to medium	Low angle from the slope, 25->28 mm	Slow push 5.5 m, slight upward tilt	Moonlit three-quarter backlight 6800 K; firelight 2100 K grazes the terrace edge only	192	8 s
S2	Below the gate lintel	Close-up	From below and beside the lintel, 28->33 mm	20 cm push, focus locked on the break	Cold light from above and the side, 7000 K; one thin highlight on the rope edge	144	6 s
S3	The slope where the rope lands	Medium	In front of and beside the fox, 40->44 mm	Slow push 1.1 m at fox eye level	Low side light 7000 K; snow bounce fills the jaw	144	6 s
S4	The slope outside the gate	Wide tracking shot	From the side of the slope, 36->41 mm	Track 6 m, speed matched to the rolling bell	Moonlit three-quarter backlight 7000 K; one cold highlight on the bell	144	6 s
S5	The slope up to the snow spine	Wide	From behind and to the side, 44->50 mm	Track 7 m with a slight rise	High cold light 7500 K; wind-blown snow scatters it into haze	144	6 s
S6	The gate below the snow spine	Wide to medium	Low wide angle, 28->33 mm	Slow push 1.4 m as the wings press down from the top of frame	Cold light from above and behind, 7500 K; under the wings almost black, silhouette only	144	6 s
S7	Shelter under the rock wall	Medium	From outside the wall, 38->43 mm	Slow push 1.2 m, settling on the triangle of three heads	Cold light from front and side, 7000 K; faint bounce off the rock wall	144	6 s
S8	Top of the snow spine	Extreme wide	Straight on and slightly low, 31->34 mm	Slow push 2 m, symmetrical composition	Pre-dawn daylight 7500 K; one colder blue line under the cloud break	192	8 s
S9	Between the spine and the rock terrace	Wide	From behind and to the side, 29->32 mm	Track and rise 4 m	High-altitude cold light 8000 K; a faint warmth in the east has not reached the snow	168	7 s

#	Length	Shot size	Camera / lens	Movement	Key light and colour temperature	Frames	Duration
S10	The rock terrace in front of the gate	Medium close	Low angle, 34->39 mm	Slight pull-back 1.2 m, bell and three animals in one frame	Cold ambient 8000 K; warm gold 2400 K at the bell core is the only moving light	144	6 s
S11	The gate and the whole valley	Wide to extreme wide	High angle, 27->30 mm	Pull back 9 m and rise 5 m	Low warm sun 3000 K backlight; the environment stays cold blue	168	7 s
S12	The valley, wide	Extreme wide	From the slope below, 26->29 mm	Slow push 3 m, then hold still for 24 frames	Low warm sun 3000 K backlight; snow reads gold #FFD9A8	168	7 s

Continuity anchors (copy them next to the shot list): The chip in the bell mouth stays front-right throughout; the broken rope end points up in S2, and the end the fox carries from S5 is the same rope; the llama's saddle blanket, packs and empty bell slot never swap sides across the 12 shots; the condor comes from above the lintel in S6 and perches top-left of the same wall in S7; the three silhouettes in S12 must match the blocking in S8. Writing these five down beats writing keep it consistent in a prompt.

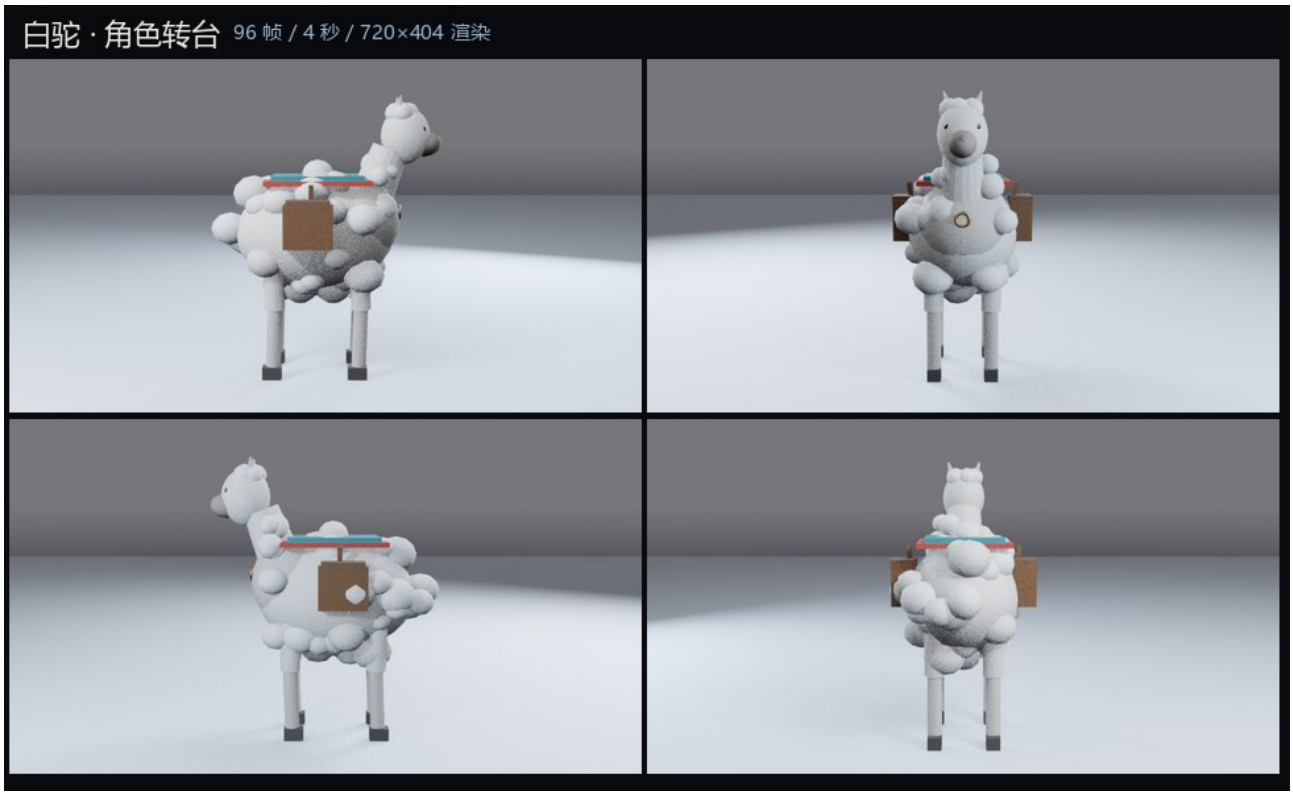
Character, prop, set: seven subject cards lock the look

Having fewer subjects makes this film easier to break - the audience has nothing to look at but three animals and one bell, so any drift in fur colour, build or bell proportion is instantly visible. Each card therefore states only three things: **scale, material, forbidden drift**. Below are the seven cards as used.

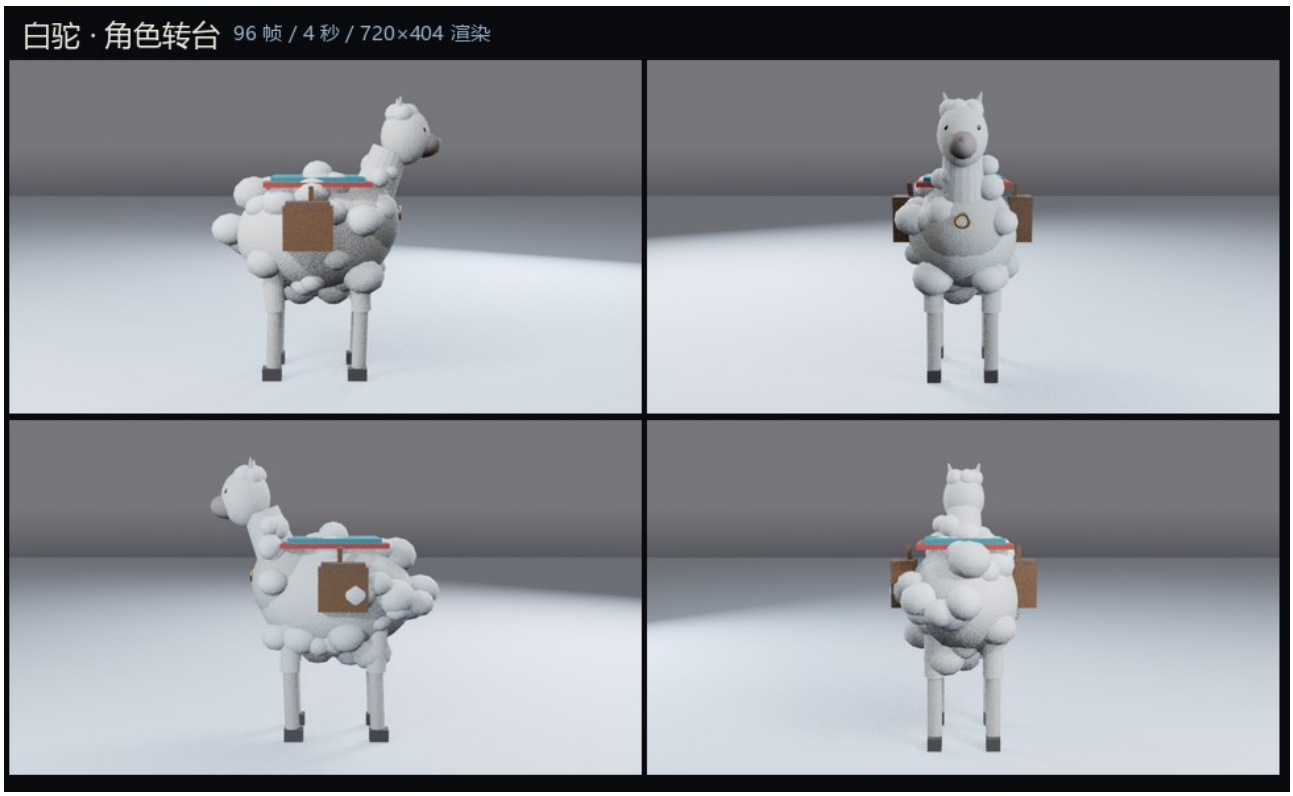
The seven subject cards and their locked items

Subject	Type	Scale and structure	Material and colour	Forbidden drift
White llama	Character	Shoulder height 1.15 m, neck 0.7 m, long slim legs; saddle blanket 1.1 x 0.9 m, two packs, empty bell slot on the right	Cream-white short coat, slightly soiled along the neck; deep indigo blanket with faded red trim; packs in raw hemp	Empty bell slot always on the right; wear on the blanket edge never moves; never becomes a camel or a horse

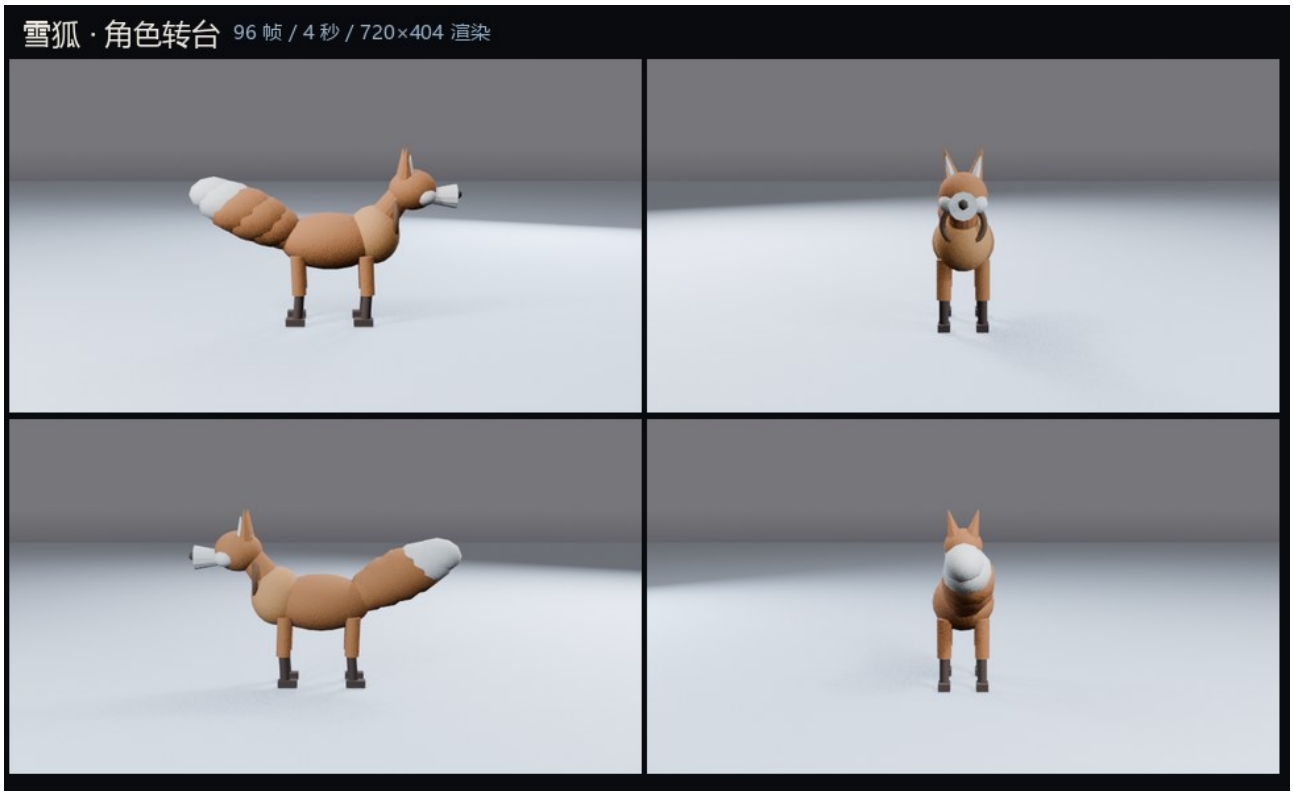
Subject	Type	Scale and structure	Material and colour	Forbidden drift
Snow fox	Character	Body 0.65 m, tail 0.45 m, rounded ear tips; about 0.38 m at the shoulder when standing	Grey-white winter coat, one pale grey stripe down the spine, pure white tail tip, near-black nose and eye rims	Tail tip must stay pure white; ears never sharpen into cat ears; never the red-brown of a red fox
Andean condor	Character	Wingspan 2.8 m, body 1.1 m; clearly splayed primary tips and a ring of white down at the neck base	Upper wing near-black with a blue-grey cast, white neck ring, dark red-grey head	Wingspan ratio never changes; the white neck ring never disappears; never becomes an eagle or a vulture
Old bronze pack bell	Prop	18 cm tall, mouth 10.5 cm across; dark iron clapper, rope threaded through the crown	Old bronze with green patina and fine scratches; the chip in the mouth sits front-right and is about 1.2 cm wide	Chip position fixed; patina never turns bright copper; never glows or throws sparks
Snow-ridge stone gate	Set	Opening 6 m high and 3.2 m wide; lintel 0.9 m thick with an old rope groove on the beam	Grey-black schist, horizontal bedding, snow only on upward faces	Opening proportions fixed; groove sits left of centre on the lintel; no carved ornament
Rock terrace and slope	Set	Terrace 7 x 5 m, slope 18 degrees; two grooves dragged into the snow by the bell	Fine-grained snow, blue-grey in shadow; exposed rock near black	Slope angle fixed; groove positions match S4 and S9
The valley and the snow spine	Set	Spine about 3 m wide, like a knife edge; valley floor about 60 m below	Snow brighter along the spine line, valley floor pressed dark; one cold blue line of sky kept open	Spine direction fixed; the valley shape in S8 and S12 must line up



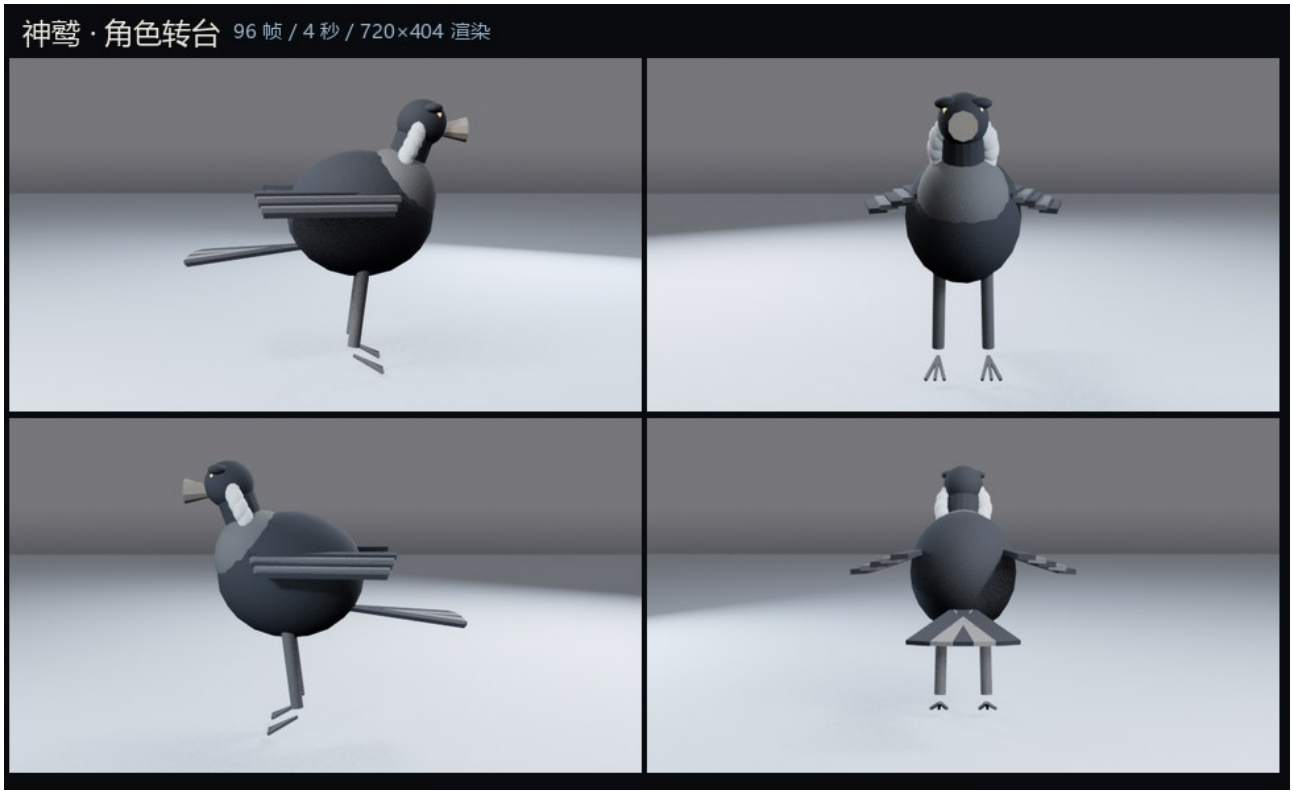
Four frames of the white llama turnaround. What matters is not whether it looks good but whether shoulder height, neck length, blanket and empty bell slot agree from all four angles. A subject that fails its turnaround does not enter final rendering.



The full 96-frame llama turnaround. The hard part of an animal character is the proportion of limbs to neck - a still hides it, a full turn does not.



Four frames of the snow fox turnaround. Check that the tail tip is pure white, the ears stay rounded and the pale grey spine stripe runs unbroken. These three are what get lost first in low light.



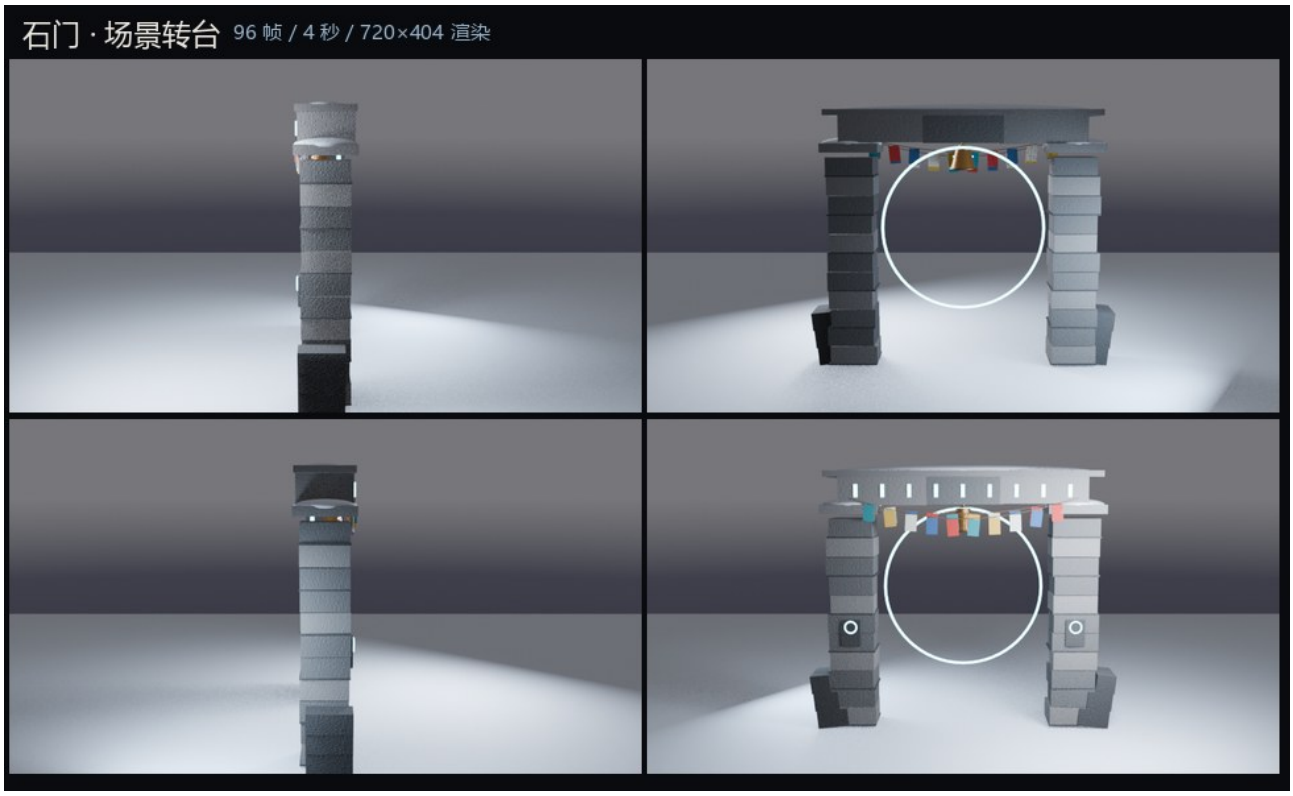
Four frames of the condor turnaround. Wingspan ratio and the white neck ring are what identify it; S6 gives only a pressing silhouette, which is exactly why the silhouette has to be locked first.



Four frames of the bell turnaround. This is the prop that must not drift: the chip in the mouth, the gap between clapper and wall, the rope hole - S2, S4, S10 and S11 all have to agree.



The full 96-frame bell turnaround. Props drift shot to shot more easily than characters do, and a turnaround is the cheapest way to pin one down.



Four frames of the gate turnaround. Check the opening's height-to-width ratio, lintel thickness and groove position. S1, S6, S10 and S11 all sit near this gate, and one distortion collapses the space.

Six collections and layered independent rendering

Each shot has its own file, and every object is filed into one of six collections: **01_character / 02_prop / 03_set / 04_light / 05_camera / 06_particles**. The point is not tidiness but that character, prop and set can render separately and be stacked back in the same frame. Anything carried in a hand or the mouth is still parented to the rig, but its collection never changes - it belongs to props.

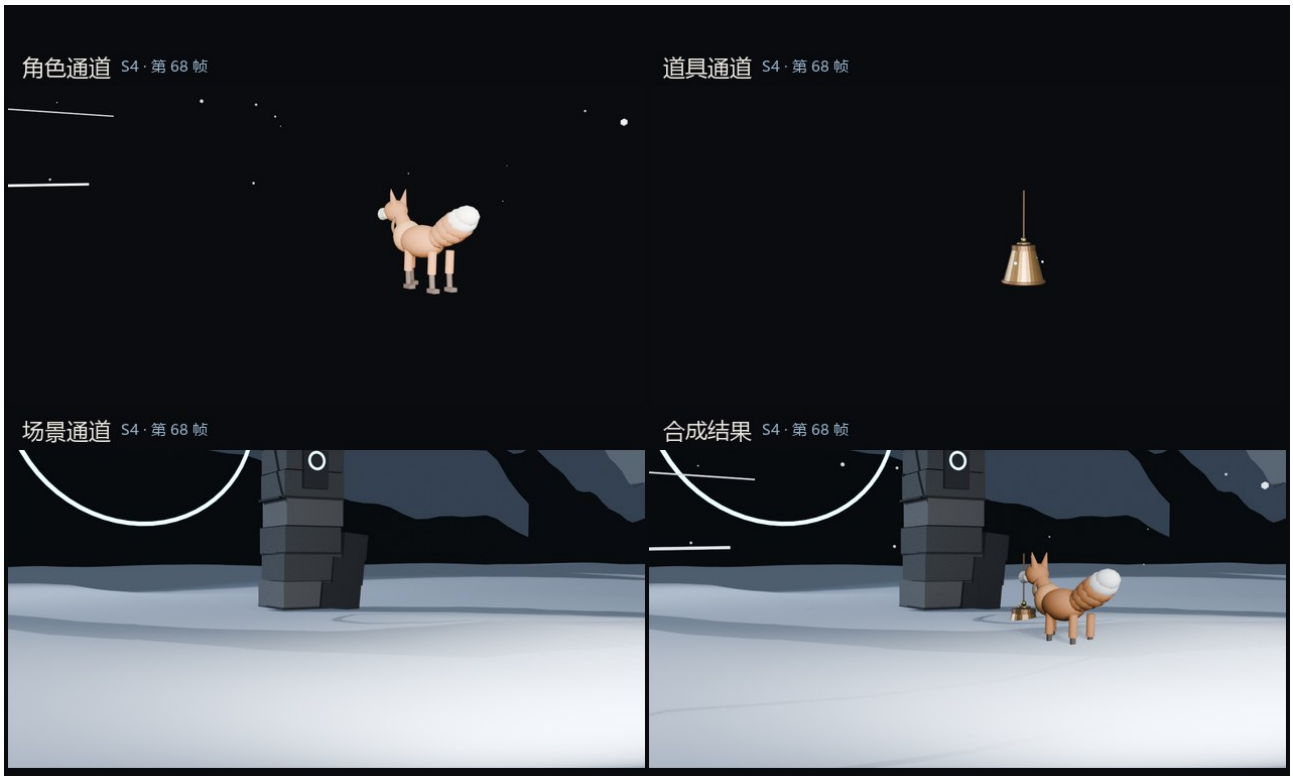
Collection boundaries and split rules

Collection	Contains	Never contains	Why render it separately
01_character	Rigs, skin and fur for the llama, fox and condor	The rope or bell in a character's mouth	Swapping backgrounds, grading characters alone, building character mattes
02_prop	Bell, rope, saddle blanket, packs	Character bodies, set rocks	Lifting the bell's warm light without polluting the environment
03_set	Gate, terrace, slope, valley, spine	Characters and props	A clean background plate for atmosphere and depth of field

Collection	Contains	Never contains	Why render it separately
04_light	Moonlight, firelight, the bell-core point light, ambient sky	Any geometry	A light pass for checking direction and occlusion
05_camera	The single camera and its focal animation for this shot	A spare second camera	Makes camera and lens reproducible across 12 shots
06_particles	Wind-blown snow, snow grains, the bell halo	Character fur	Snow and halo as separate passes, easy to darken or lift



The six-collection element matrix. Object types are on the left, their collection and render pass on the right. Its job is to stop on-the-spot decisions about where something goes - once an object changes collection between shots, the layered passes no longer line up.



S4 layer comparison: character pass, prop pass, set pass and the composite. Check three things - where the bell meets the snow, how deep the fox's paws sink, and whether the bell's shadow agrees with the moonlight. If any one is wrong, fix the shot file rather than the composite.



S11 layer comparison. This shot adds the halo and the wind-blown snow: the halo renders alone so its brightness and spread are controllable later, and the snow renders alone so it can be pushed down at the instant of the ring, letting the audience hear it.



Layer-toggle demo: character, prop, set and particles switched on one by one. Acceptance is that the remaining layers still hold with any one switched off - if the character floats once the set is gone, the contact points are wrong.

3D previz: solve space before texture

Previz answers three questions only: is the spatial relationship right, can the camera read the action, and does the key light separate the subject from the background. Material, fur and snow detail can all come later; a wrong composition has to be redone. This film's previz shares camera, lens and frame count with the finished shots, so no shot duration had to change afterwards.



The S8 composition check. Three subjects stand apart on the spine, shadows point the same way, and the cloud break stays open above. This is the film's poster shot: if it does not hold, the film loses its visual memory.

AI refinement: unify the texture, do not replace the shots

Previz solves space, refinement solves texture. This film keeps the previz camera, movement and cuts, and changes only three things on the image: **unify material and grain, press the cold palette into one hue, let the bell's warm gold stand alone**. The method is one refined still per shot as a style anchor, then the same colour settings applied across the timeline; there is no per-shot repaint, which is how one llama would otherwise become twelve.

【Project】87-second snow-epic short, 16:9, photoreal cinematic, low-saturation cold blue-grey, light film grain, no text, no subtitles, no watermark, no logo.
【Character】White llama: shoulder height 1.15 m, cream-white short coat, slightly soiled neck, deep indigo saddle blanket with faded red trim, two raw hemp packs, empty bell slot on the right; never change species.
【Character】Snow fox: body 0.65 m, grey-white winter coat, one pale grey stripe down the spine, pure white tail tip, near-black nose and eye rims; never red-brown.
【Character】Andean condor: wingspan 2.8 m, upper wing near-black with a blue-grey cast, a ring of white down at the neck base, dark red-grey head; wingspan ratio fixed.
【Prop】Old bronze pack bell: 18 cm tall, mouth 10.5 cm across, green patina, chip in the mouth front-right about 1.2 cm wide, dark iron clapper, rope through the crown; no glow, no sparks.

【Set】Snow-ridge stone gate: opening 6 m high and 3.2 m wide, lintel 0.9 m thick with an old rope groove left of centre; grey-black schist with horizontal bedding, snow only on upward faces.
【Shot】S10 medium close, 34->39 mm: the llama pulls the bell to the terrace edge, the fox braces the bell with its head, the condor lifts the rope onto the beam with its beak; once the rope holds, a ring of warm gold lights at the bell core and spreads along the mouth.
【Movement】Low camera, slight 1.2 m pull-back, even speed, hold at the end; no pan, no crane, no orbit.
【Consistency】Chip on the same side; broken rope end points up; blanket and empty slot never swap sides; blocking matches S8; one main action per shot.
【Negative】Species swap, warm fur, bright copper or crystal bell, chip on the other side, shorter wingspan, blown-out firelight, camera shake, slow motion, extra people.

1. **Set one style anchor first:** Pick the shot that best represents the texture (here S8 on the spine), refine it until you are satisfied, then write that colour, grain and contrast into settings and apply them to the other shots.

2. **Change one thing per shot:** One shot solves one problem: material, or light, or snow surface. Change three at once and you cannot tell which one broke it.

3. **Re-check cuts after refining:** Refinement changes brightness, so brightness steps at the cuts have to be checked again. Keep the snow brightness difference between adjacent shots under 15 percent; beyond that, add a transition layer.

4. **Keep the previz version:** Save refinement as a new version instead of overwriting previz. When you rework a shot you need the spatially correct version back, not a rebuild from zero.



The refined poster frame: gate, white llama, snow fox and a touch of warm gold on the bell. Used for the tutorial cover and the title card, not in the main cut.



Refined bell still. Where the chip sits, how the patina spreads and the gap at the clapper are the three things this frame hands to every later shot.



Refined S1 still. Moonlit backlight presses the gate into silhouette and leaves the llama a single rim; the empty terrace on the right is the same spot S10 will use to hang the bell.



Refined S9 still. The rope is a straight line and the bell ploughs a deep groove. The effort in this shot lives entirely in that line and that groove; refinement must not smooth them away.

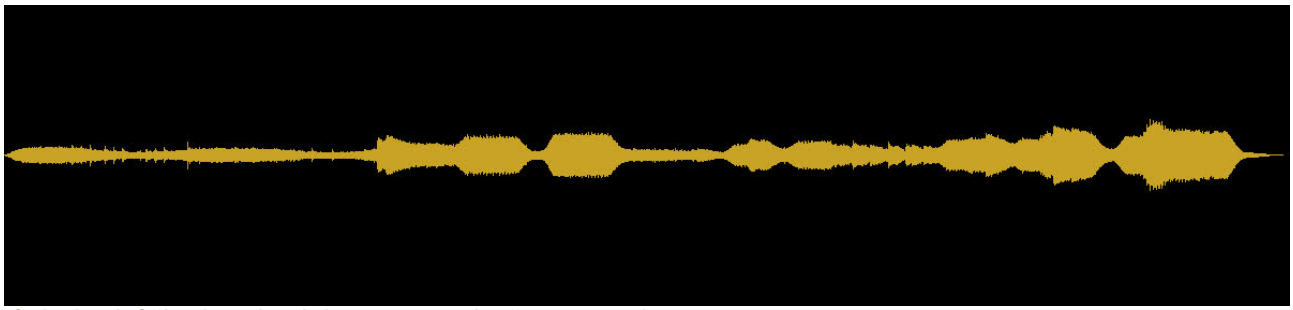


Refined S12 still. Snow starts to turn gold while the sky keeps one cold blue line. This is the only place the film lets warmth fill the frame; earlier and dawn loses its meaning.

Four sound layers and three rings

Sound is built in four layers from the bottom up: **ambience** (wind through the gate, snow grains on stone, wind pressure at altitude), **action** (rope fibres tearing, hooves sinking, the bell scraping snow, rope landing on the beam), **emotion** (three breaths at different depths, the fox's short sniffs), **transitions** (the half-beat where wind stops, the bell's decay). There is no melodic score - the film's temperament depends on emptiness, and music would turn it into a different film.

The three rings are the film's only loud events, and their positions are fixed: **the first after the bell is hung in S10**, close, dry, bell chamber only; **the second as the halo clears the gate in S11**, with valley echo; **the third as S12 closes**, low frequency and tail only, falling into silence. Leave enough gap between them, or the bell starts to sound like an alarm.



If playback fails, download the 87-second mix (AAC, 48 kHz, stereo, 1.7 MB).

The 87-second four-layer mix. Listen to ambience alone first and confirm the wind direction and strength change at S5 and S11; then action alone, confirming every hoof falls on picture; then all three rings together.

Platform film: one timeline, 14 video clips plus one audio track

The film does not use auto assembly; it is built by hand in the platform's film editor: **title card 0-4 s -> S1 4-12 -> S2 12-18 -> S3 18-24 -> S4 24-30 -> S5 30-36 -> S6 36-42 -> S7 42-48 -> S8 48-56 -> S9 56-63 -> S10 63-69 -> S11 69-76 -> S12 76-83 -> end card 83-87**, plus one audio track from 0 to 87 s. That is 14 video clips and one audio clip, 87 seconds in total.

87-second delivery timeline

Segment	In	Out	Content	Duration
C0	00:00	00:04	Title card	4 s
S1	00:04	00:12	Gate and white llama	8 s
S2	00:12	00:18	Rope snapping, close-up	6 s
S3	00:18	00:24	Snow fox looks back	6 s
S4	00:24	00:30	Bell rolls down	6 s
S5	00:30	00:36	Into the wind, uphill	6 s
S6	00:36	00:42	Condor bears down	6 s
S7	00:42	00:48	Three animals in shelter	6 s
S8	00:48	00:56	Standing apart on the spine	8 s
S9	00:56	01:03	Hauling the bell uphill	7 s
S10	01:03	01:09	Hanging the bell at the terrace	6 s
S11	01:09	01:16	The ring returns to the mountains	7 s
S12	01:16	01:23	Dawn close	7 s

Segment	In	Out	Content	Duration
C12	01:23	01:27	End card	4 s

Editing does three things: order by the table, trim to the action, breathe at the two black gaps. One black gap goes from title into main film, one before the end card; everything else is a hard cut landing three frames before the action's wind-up. S11 is never cut, so the audience travels with the halo from the gate across the whole valley.

Acceptance: let the files speak

Measured parameters of the delivered files

File	Resolution	Duration	Frame rate	Audio	Size
720p master	1280×720	87.000 s	24 fps	AAC, 48 kHz, stereo	39.2 MB
Platform 480p export	854×480	87.000 s	24 fps	AAC, 48 kHz, stereo	5.0 MB
Standard 480p	854×480	87.000 s	24 fps	AAC, 48 kHz, stereo	6.9 MB

Completion criteria: All 12 shots play in sequence; the llama's blanket and empty bell slot, the fox's tail tip, the condor's wingspan and the bell's chip hold across shots; the bell meets the snow convincingly in S4; the warm gold at the bell core in S10 is the only moving light in frame; all three rings land after 01:03 and never overlap; the platform 480p and the local 720p master both stream from a direct link at 87.000 seconds, 24 fps, with a stereo track, and each of the three rings falls inside the shot the shot list assigns it to. Uploaded assets, a single candidate or a queued export are not completion.

Do it yourself: rebuild this film in your workspace

1. **Create the blank canvas and fix its boundaries:** Create a blank canvas and fix the boundaries: title, 87 seconds, 16:9, 12 shots, no dialogue. Build seven subjects first - white llama, snow fox, condor, bell, gate, terrace and slope, valley and spine - each with one main image and up to two supporting images.

2. **Write the five acts and the 12-shot table:** Set the five-act structure and the shot list first, then fill in shot size, camera position, focal length, movement, lighting and frame count per shot. Bind each shot to its asset cards (characters, props, scene state); the whole film should land near 53 bindings. Write one main action per shot. Appearance lives in the cards; action lives in the shot description.

3. **Build 3D previz:** Block the low-poly set, then fix camera, lens and movement before light positions. Render first, middle and last frames per shot, checking for black frames or jumps and keeping brightness variation within one shot under 15 percent.

4. **Split the file into six collections:** One file per shot, objects filed into 01_character / 02_prop / 03_set / 04_light / 05_camera / 06_particles. A rope held in the mouth is parented to the rig but stays in props.

5. **Render four passes:** Render character, prop, set and a white-model pass separately, keeping light and camera visible in every layer. Composite three layers back on one frame and check edges, contact points and cast shadows.

6. **Pass the turnarounds first:** Give the llama, fox, condor, bell and gate 96-frame turnarounds each. A subject that fails its turnaround does not enter final rendering.

7. **Refine for one texture:** Make one style anchor, then change one thing per shot; re-check the brightness step at each cut and keep the previz version for rework.

8. **Edit and lay the sound:** Order the 14 segments, use black gaps only at the head and tail, hard-cut everything else. Lay ambience first, then action, then the three rings; never start with music.

9. **Export and check every field:** Export 480p through Manual film, then produce the 720p master from the same sequence. Check resolution, duration, frame count, audio track and file size, then actually play both files.

Asset and cost boundaries

□ Inside the platform: Seven subject cards, five turnaround sets of 96 frames each, a 12-shot contact sheet, the six-collection matrix, two layer comparisons plus the layer-toggle demo, 13 refined stills, the 87-second master and the platform 480p export - all published with this tutorial for direct comparison.

□ Re-creation: Uploading, canvas organisation, storyboard structure, manual film, editing and export are available; generation models are billed per job, so check the quote before confirming and do not submit paid jobs that were not authorised.

□ Remix: You can reshoot your own version from the shot list; if you reuse this project's characters and sets, keep the attribution so readers do not mistake them for your own designs.

For canvas basics go back to Case 2; to practise layered work with several characters read Case 5; to train shot language properly start with Advanced creation.

Case 7: Complete dual-line narrative film • From script to "Cloud Ridge Letter"

What you see first: An 81-second, 12-shot Cloud Sea Postal Route short film, "Cloud Ridge Letter", ships with a 720p master, a platform 480p export, a 12-shot contact sheet, an 81-second timeline, 5 turntable groups, and 18 retouched stills. The film has no dialogue; it holds together by showing two people in two places doing the same thing at the same time: **Lan guards a letter in the clouds, while the old station master lights a backup lamp in the lighthouse.** Case 6 practiced "few characters, strong light"; this one practices something harder-**Dual lines** 11 of the 12 shots follow the protagonist; only S8 cuts back to the shore. This cut back is not a flashback but the stitch that sews the two lines together.

Studio → episodic project → script and five acts → script breakdown → 30 asset cards → 12 shot nodes → 3D previz → layered render → dual-line edit → AI finish → four-layer sound → export

Live project: This film is saved in the studio as a formal project: the workflow project " Letters from the Cloud Ridge" holds the 81-second spec, the 38-paragraph script and the 12-shot list; the breakdown is confirmed through revision 2 (25,219 characters, 30 asset requirements), and all 30 cards in the asset library are locked. The canvas "Letters from the Cloud Ridge · 12-shot production canvas" has **111 nodes and 82 links**: 81 media nodes (8 process notes, 12 storyboard nodes, 12 previz stills, 12 single-shot previz clips, 5 character/prop references, 5 turntable sets, 3 layered comparisons, 18 finished stills, 2 title and end-card clips, 2 master assets and 2 film nodes) plus 30 asset-card nodes. The links are 18 media references, 11 parent-to-state cards and 53 shot bindings. Open it for comparison after signing in: Cloud Ridge Letter · 12-shot production canvas. The film was edited by hand in the platform and exported as one video track with 14 segments plus one audio track.

See the result first: one master, one platform export



If it will not play, download the 720p master (1280×720, 81.004 seconds, 24 fps, AAC 96 kHz stereo, 53.5 MB).

The 81-second delivery version of "Cloud Ridge Letter". The structure is a 4-second opening title card + 73-second main film (12 shots, 1752 frames) + 4-second closing title card, 1944 frames total. Watch it twice: the first pass only look at "who has the letter"; the second pass only watch the lamp-you'll see the lamp goes out at S4, relights at S8, and shines on a face at S11.

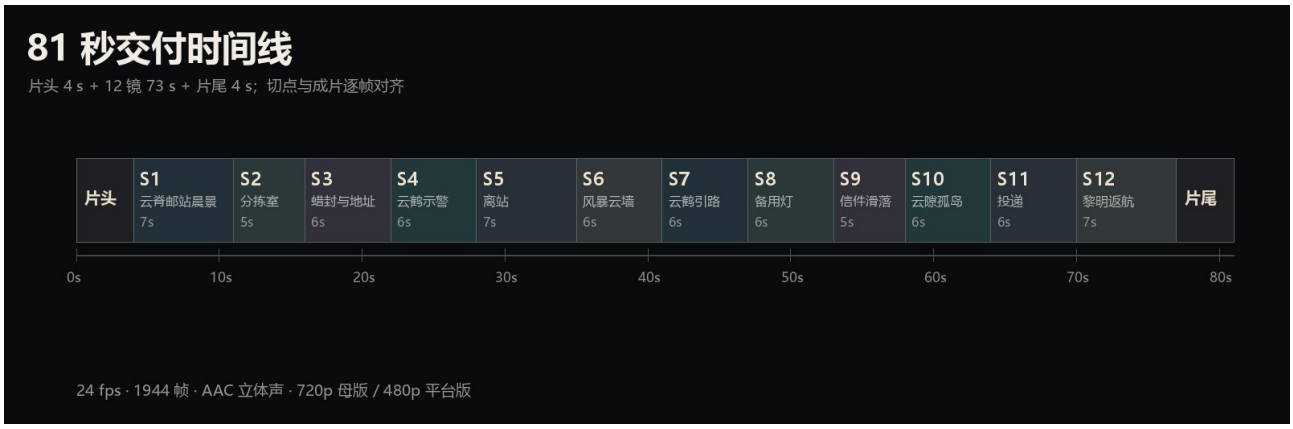


If it will not play, download the platform 480p export (854×480, 81.0015 seconds, 10.0 MB).

This file is what you get by putting the same 81-second timeline through the platform's "Manual film → Export" path: 854×480, 81.0015 seconds, 24 fps, AAC 96 kHz stereo, 10.0 MB. Both files are 96 kHz stereo and cut at exactly the same points; only resolution and bitrate differ. Acceptance always reads the values ffprobe reports, never the file name or the thumbnail.



Contact sheet of the 12 previsualization shots, in final film order. This image is a "construction drawing," not a " beauty render": it only answers whether spatial relationships, camera positions, and focal lengths are correct. Check three things-whether the windowsill in S3 and the cloud wall outside are on the same line of sight, whether the dock in S5 is really at the edge of the stone ridge, and whether the isolated island in S10 falls on the golden section point of the frame.



81-second timeline. The opening and closing are each 4 seconds; among the 12 shots, S1, S5, and S12 are each 7 seconds, S2 and S9 are each 5 seconds, and the remaining 7 shots are each 6 seconds, totaling 1752 frames in the main section. S9 is only 5 seconds but placed at the narrowest point of the film-it is the only shot where "the character might lose the mission," and a shorter length keeps it from dragging.

One-sentence story and the five-act causal chain

In one sentence: **Above the sea of clouds, on the last mail route, a young courier takes a letter pressed with red wax, flies through the storm cloud wall, and drops it into the old mailbox on the isolated island at the moment the lighthouse shines again.** no dialogue, so causality must be written in actions: **Receives the letter → light goes out → still takes off → enters the cloud wall → crane guides → old man lights the lamp → letter slips from hand → catches it → sees the island through a cloud gap → delivers the letter → light falls on her face → returns**. Every link is a visible action; the audience doesn't need to be told.

Five-act structure and causal chain (consistent with the 81-second final cut)

Act	Timecode	Function	Key action	emotion
0 Title	00:00-00:04	Set the tone, give the title	Black screen; the title "The Cloud Ridge Letter" and subtitle fade in; wind rises.	Still
1 Receives the letter	00:04-00:22	Establish mission and omen	Airship arrives at the station; the old man takes the last cross-cloud letter from the drawer; the wax seal pattern develops a fine crack in the morning light, and the cloud wall outside presses in.	Solemn, with a faint unease.
2 Warning	00:22-00:35	First reason to refuse departure	A cloud crane skims past the ridge, and the lighthouse light dims at the same time; Lan still casts off and leaves the station, while the old man stands on the shore holding a lamp.	Alert, restrained.
3 Crossing the clouds	00:35-00:53	Force the protagonist into blindness.	The gondola tilts sideways in the cloud wall; the crane pushes the fog away at the bow; cut back to the lighthouse, where the old man turns the valve, strikes a match, and lights the backup lamp.	Pressure → a glimmer of hope
4 Guarding the letter	00:53-01:04	Cost and turning point	The letter slips out of the mailbag; Lan unbuckles and leans out to catch it, her shoulder hitting the cabin wall; the clouds split open a vertical gap, and the isolated island appears.	Brief peril → sudden expanse
5 Delivery	01:04-01:21	Pay off the promise	The letter falls into the tin mailbox with a soft sound; the lighthouse beam sweeps through the clouds onto her face; the Skylark returns, its mailbag empty.	Relaxed, answered.
6 End card	01:21-01:25	Leave an aftertaste	Warm light spreads across the sea of clouds, and the title fades out.	Still

Why it must cut away at S8: The most common mistake in a dual-line film is that the second line is just exposition. Here, every second of S8 serves what follows-viewers watch the old man turn the brass valve, strike a match, and the light first spread into a small point in the lens, so that in S11 they recognize where the light sweeping across Lan's face comes from. If S8 were omitted, the light in S11 would be just a pretty effect; with S8, it becomes " someone lit the lamp for me." **The only test for a dual storyline: if you delete the other line, would this one collapse?**

The 12-shot list is the construction drawing

The shot list contains no adjectives, only checkable parameters: shot size, camera position and focal length, movement, light position and color temperature, frame count, duration. The table below is the version actually used for this film; the main cut totals 1752 frames and 73 seconds, matching the finished film frame by frame; plus 96 frames each for opening and closing titles, the total is 1944 frames and 81 seconds.

"Cloud Ridge Letters" 12-Shot Production Sheet (main film 73 seconds)

#	Length	Shot size	Camera / lens	Movement	Key light and colour temperature	Frames	Duration
S1	Cloud Ridge Post Station · Morning Exterior	Extreme wide to wide	Low camera on the ridge spine centerline, 30 mm	Truck and gentle push in; start frame on the cloud sea, end frame with the two at the lower-left third	Morning side-backlight 4300 K; window light 2700 K only grazes the porch	168	7 s
S2	Sorting Room · Morning Interior	Over-the-shoulder two-person close-up	Slightly below eye level, 27 mm	Breathing slight drift; foreground is bundles of letters	Window light 8000 K as key; oil lamp 2700 K fills the jaw	120	5 s
S3	Sorting Room by the Window · Morning Interior	Close-up to macro	Nearly touching the table, 62 mm	Focus pulls slowly from the wax seal to the cloud wall outside the window	Window light 9000 K hard side; warm fill 3200 K only reaches the paper edge	144	6 s

#	Length	Shot size	Camera / lens	Movement	Key light and colour temperature	Frames	Duration
S4	Stone Ridge East Side · Morning Exterior	Wide tracking shot	Low on the ridge side, 40 mm	Gently pan with the crane; the crane enters right and exits left; the lighthouse stays in focus in the background	Cool skylight 9000 K; lighthouse 2700 K warm light is the only warm source in the frame, then goes out	144	6 s
S5	Mooring Mast · Morning Exterior	Full shot to low-angle shot	Close to the dock, 26 mm	A small tilt-up following the pod's rise	Morning side-backlight 4200 K; cool fill 9000 K holds the pod underside in shadow	168	7 s
S6	Inside the Cloud Wall · Noon Exterior	Wide	Mounted outside the pod, 44 mm	Follows the ship's roll; the horizon is kept at a 12° tilt without correction	Storm 9500 K top-side; searchlight 4300 K cuts a white streak	144	6 s
S7	Deep in the Cloud Wall · Noon Exterior	Medium	Bow, 36 mm	Push in with the crane; the crane stays slightly above center in the frame	Fog scatter 9500 K; the crane's white forehead feathers are the only bright spot	144	6 s
S8	Lighthouse Interior · Noon Interior	Medium close	High angle down from the iron ladder, 39 mm	Fixed; when the lamp lights, exposure shifts from outside the window to the face	Skylight 9500 K accounts for only 12%; lamp core 2200 K spreads from a point to fill the lens	144	6 s

#	Length	Shot size	Camera / lens	Movement	Key light and colour temperature	Frames	Duration
S9	Skylark Pod Interior · Afternoon Interior	Medium	High inside the pod looking down, 35 mm	Tilts with the ship; the envelope stays at the center of the frame	Cabin window light 9000 K as key; lantern 2700 K gives the wooden boards a warm edge	120	5 s
S10	Cloud Gap · Afternoon Exterior	Extreme wide	From the ship's side looking downward, 40 mm	Slow descent follow; the solitary island stays on the golden section point	Cloud gap light shaft 4300 K; surroundings are pushed down to 9000 K dark blue	144	6 s
S11	In Front of the Old Mailbox on the Solitary Island · Dusk Exterior	Medium close	At the same height as the mailbox, 48 mm	Fixed; the lighthouse beam sweeps forward from the right rear	Dusk 4200 K back-side; lighthouse beam 2700 K sweeps across the side of the face	144	6 s
S12	Above Cloud Ridge · Dawn Exterior	Extreme wide shot	From ahead of the bow looking back, 33 mm	Slow pull-back to bring the airship, crane, and post station into the frame together	Low-angle dawn 4500 K floods the cloud surface; post station lights form strings	168	7 s

Continuity anchors (copy them next to the shot list):① The gap in the cloud-ridge pattern on the wax seal is always right of center on the front; the thin crack in S3 and the one dropped into the mailbox in S11 are the same crack. ② Lan's old copper hair clip is behind her left ear; the cloak's ochre yellow is the second warmest color in the cool palette, only one step darker than the lamp. ③ The old stationmaster's lantern is carried in S1 and S5; S8 switches to the backup lamp in the tower-these two lamps are not the same one. This must be visibly communicated, otherwise S8 is wasted. ④ After the lighthouse lamp goes out in S4, no light may appear in the cloud gaps in S10; only S11 may show it lit. ⑤ In S12, Lan's mail bag must be empty. Locking down these five rules is far more useful than writing "stay consistent" in the prompt.

Script breakdown → 30 asset cards: characters, props, scenes and time of day

These 30 cards are not a hand-written list; they come out of the script breakdown. The episode project holds 38 script paragraphs first (P1-P38), and the breakdown turns every noun that will certainly be drawn into a requirement: **3 characters, 5 props, 11 scenes, 11 scene states**. Each requirement carries its original line and fact sheet into the asset library; you fill in scale, material and drift guards one by one, upload one primary reference, then lock it. Once locked, that card is the only basis for shot binding - the 12 shots carry 53 bindings, so changing one card changes the whole film. The table below covers the **seven lead cards** of these 30 that are refined first; turntable measurements are in parentheses, and the full set of 30 is listed item by item further down.

This chain can be opened row by row on the platform: Every requirement on the breakdown page carries its current card state right after the name - "card pending", " locked / primary reference" or "reference missing". The asset-card chain table below puts **requirement -> source paragraph -> asset card -> version -> primary reference** in six columns, with four counters above it: active requirements, cards created, versions locked, primary references present. All four read 30 for this film, which means every requirement the breakdown raised landed in the asset library; if any number disagrees, the table shows exactly which row broke. This table is also the only hand-off receipt between the script breakdown and the assets stage in this course: a requirement is not pasted into a document, it points at a card that can be locked, reused and bound to shots.

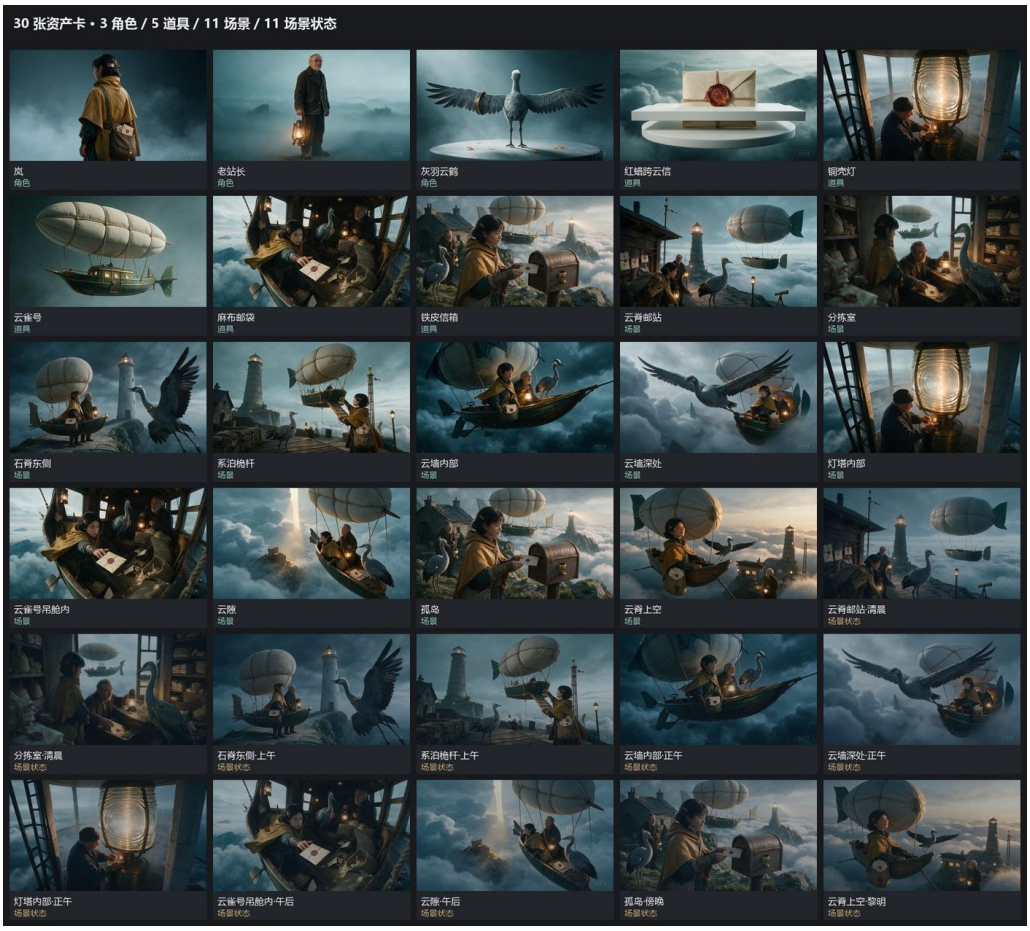
Seven lead cards inside the 30, finished first (turntable measurements in brackets)

Subject	Type	Scale and structure	Material and colour	Forbidden drift
Lan	Character	Around twenty, about 1.70 m tall; short black hair pinned behind the left ear with an old copper clip, wearing a mail bag across the body (turntable bounding box 0.72×0.60×1.63 m, 34 pieces, 2,280 triangles).	Ochre yellow waxed-cloth cloak (with rain stains and a worn hem), dark blue high collar, leather wrist guards, coarse linen mail bag.	Hair clip behind the left ear; cloak is not to be changed to bright yellow or orange-red; do not change face shape or height.

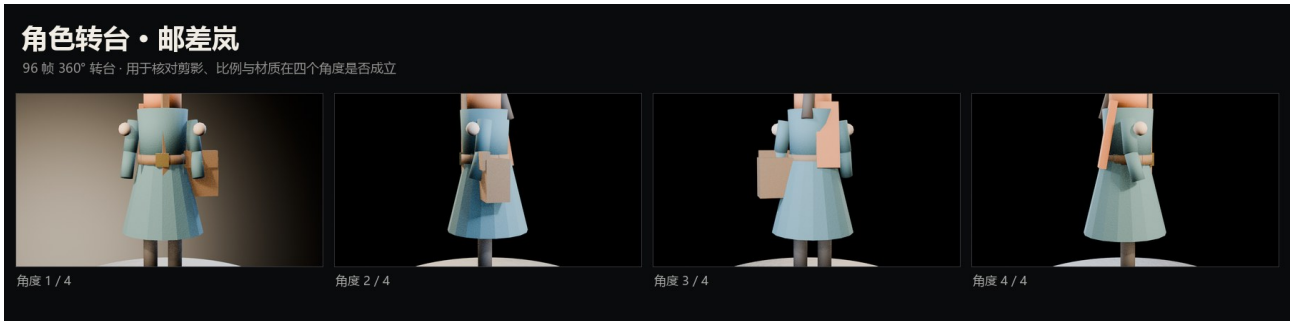
Subject	Type	Scale and structure	Material and colour	Forbidden drift
Old Stationmaster	Character	In his sixties, about 1.58 m tall, slightly stooped; gray-white short beard, round copper-rimmed glasses (turntable 0.61×0.62×1.58 m, 23 pieces, 2,160 triangles).	Dark brown wool coat (shiny at the elbows), gray-white short beard, copper-cased lantern.	Glasses frame must be round copper-rimmed; the lantern is not the same lamp as the backup lamp in S8; do not make him younger or change hair color.
Gray-feathered cloud crane	Character	Standing height 2.2 m, wingspan set to 4.2 m; turntable freeze-frame in gliding pose, bounding box 3.63×2.03×1.74 m (71 pieces, 2,256 triangles).	Blue-gray feathers, wingtips near black, a tuft of white feathers on the forehead, a ring of old copper on the inner side of the right wing.	Wingtips must be near black; the white forehead tuft must not disappear; copper ring is on the inner right wing; do not replace with egret, eagle, or phoenix.
Skylark	Vehicle	Gasbag 12.2 m long, wooden elongated shuttle-shaped gondola; turntable bounding box 4.20×14.05×4.96 m (84 pieces, 3,484 triangles).	Dark green paint with brass trim, double-layer gasbag interlaced by three ropes, twin tail fins at the rear, wooden floor and brass handrails in the cabin.	Gasbag must be double-layer; the interweaving of the three ropes must not change; tail fins must be paired; do not change into a hot-air balloon or propeller aircraft.
Red Wax Letter	Prop	Envelope aspect ratio about 1.5:1; wax seal occupies about one third of the flap center; hemp cord wrapped once around the flap (5 pieces, 100 triangles).	Beige envelope, dark red wax seal embossed with a cloud-ridge pattern, natural hemp cord color.	Wax seal must be dark red, not true red; the pattern gap is on the front, right of center; hemp cord must not disappear; no stamps or writing added.
Cloud Ridge Post Station	Set	Stone ridge about 30 m wide and 64 m long; the post station is at the southwest end; timber-framed main building + mooring mast + drying line for letters; the yard can park one 14 m airship.	Gray-blue schist, dark brown timber frame, window lamps 2700 K; roof and mast top show years of rain staining.	Post station is always at the southwest end, lighthouse at the northeast end; the mail-drying line must not change position; no modern buildings.

Subject	Type	Scale and structure	Material and colour	Forbidden drift
White Stone Lighthouse	Set	At the northeast end of the stone ridge, white stone tower, with an iron ladder, huge lens, brass valves, and backup lamp in the lamp room.	White stone with gray mortar joints, brass valves and lens mount darkened; lamp wick color 2200 K.	The tower stays at the northeast end and does not move; after S4 goes dark, the tower top must not light again until S11; the lens must be a multi-faceted prism structure

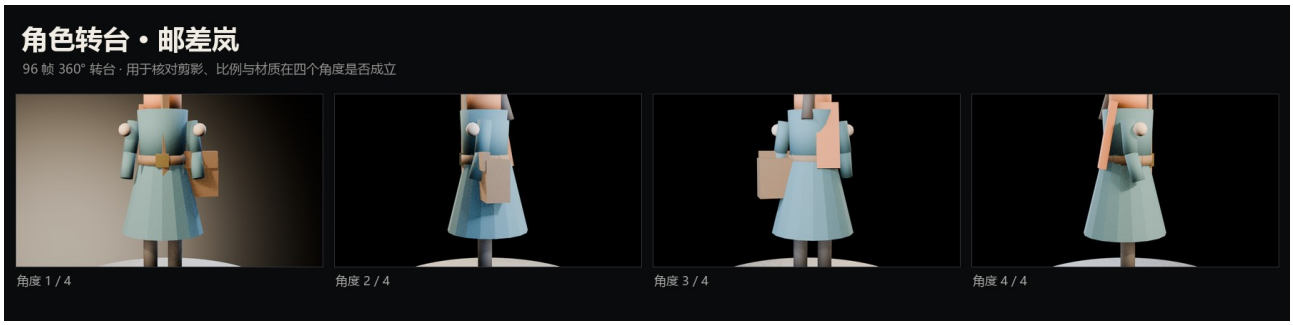
How the seven lead cards map onto the 30: Asset cards are split by what the camera will actually see, not by what exists in the world. That is why the lighthouse exterior has no card of its own - it is written into the lock clauses of "Scene | East Ridge" and "Scene | Cloud Ridge Post" as a landmark, while the interior space is the separate card "Scene | Lighthouse Interior". For the same reason the skiff is both a vehicle and an interior set, so it appears in the 5 prop cards and in the 11 scene cards: one manages its silhouette, the other manages the doors, handrails and sliding crates inside the gondola.



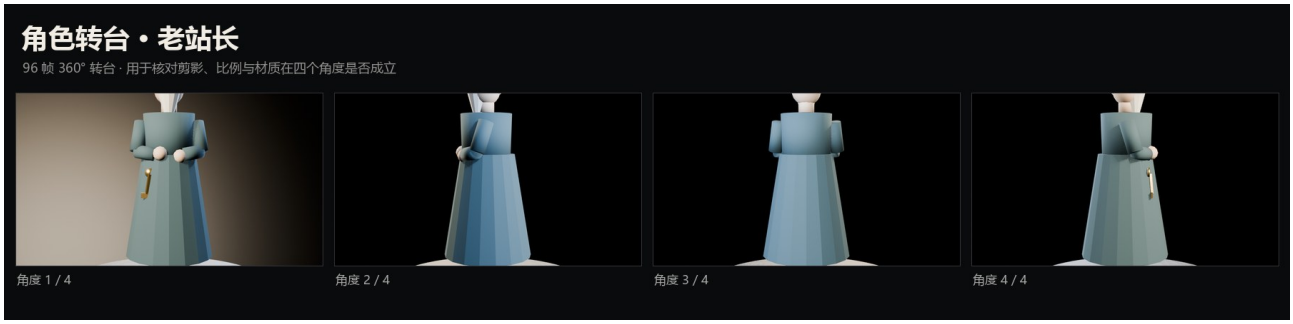
Contact sheet of the 30 asset cards, 5 columns × 6 rows, ordered 3 characters, 5 props, 11 scenes and 11 scene states. every card carries a name, a fact list, lock clauses and one primary reference - the contact sheet is only a human-readable index; the real cards live in the asset library with every status column reading "locked". Printing it and taping it beside the monitor is the cheapest habit in this workflow: copy from the card when you write prompts, do not work from memory.



Four frames from Lan's turntable. The point is not "does it look good," but whether the cloak drape, hair clip position, and cross-body mailbag height are consistent across the four angles. A subject that fails the turntable check does not enter final rendering.



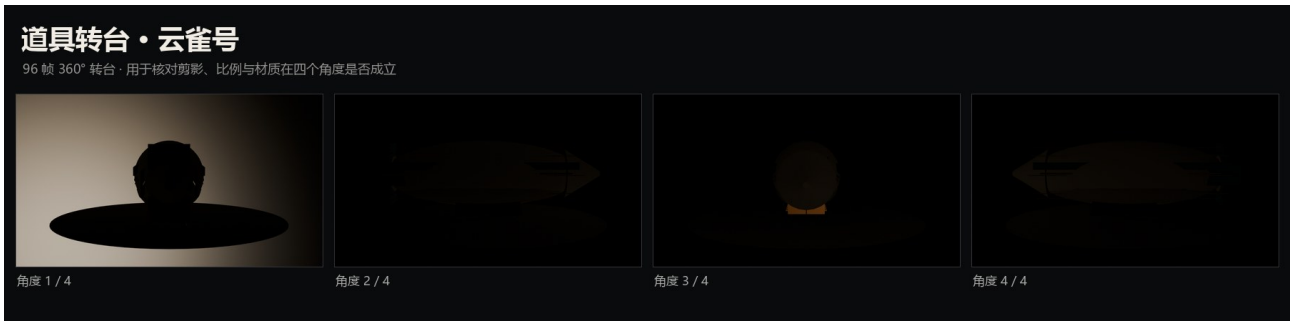
Lan's full 96-frame turntable (854×480, 4 seconds). The hard part with humanoid characters is the relationship between shoulder width and the cloak-a still frame cannot show whether the cloak will clip during a turn, but a full rotation will reveal it.



Four frames from the old station master's turntable. Check the arc of the hunched back, the round glasses, and which hand holds the lantern. In S5 he stands on the shore holding the lamp-the last frontal view in the film-so the silhouette in that frame must hold up.



Four frames from the cloud crane's turntable. The near-black wing tips, the white forehead feathers, and the copper ring inside the right wing are the three points by which it can be recognized in fog. In S7 it is the only thing lit in the whole shot; precisely because of that, these three details must be locked down first.



Four frames from the Skylark's turntable. This is the most "expensive" subject in the film: the layering of the gasbag, the interweaving of the three ropes, and the proportions of the twin tail fins-if any of these changes, the audience will immediately feel "this is not the same ship." The 96-frame turntable covers four directions; that is enough.



Four frames from the red-wax letter's turntable. Props are more likely than characters to "look different in every shot": the dark red saturation of the wax seal, the direction of the pattern notch, and the number of rope loops must match across S2, S3, S9, and S11. A turntable is a low-cost way to lock down the appearance.

The 30 cards: 3 characters / 5 props / 11 scenes / 11 time-of-day states

Naming is fixed when the breakdown is confirmed, and every node, shot binding and prompt afterwards uses that name: characters as "Character | Name", props as "Prop | Name", scene parents as "Scene | Location", scene states as "Scene state | Location · time of day". The payoff is direct - reading the shot list, you never guess which card S9 uses, and you never re-describe it inside a prompt.

3 character cards

Asset card	Requirement key	Script paragraph	Scale and structure	Lock clauses
Character | Lan	char_lan	P8	About 1.70 m tall; short black hair with an old brass clip behind the left ear; rough linen mailbag worn cross-body	The clip stays behind the left ear; the cloak is ochre waxed canvas, never bright yellow or orange-red; face and height do not drift

Asset card	Requirement key	Script para graph	Scale and structure	Lock clauses
Character | Old Stationmaster	char_stationmaster	P9	About 1.58 m tall, slightly stooped; short grey beard, round brass-rimmed glasses; dark brown wool coat	The glasses must be round with brass rims; the hand lamp and the backup lamp in the tower are two separate props and must never merge into one
Character | Grey-Feather Cloud Crane	char_crane	P10	2.2 m standing height, 4.2 m wingspan; blue-grey plumage, near-black wingtips, a tuft of white feathers on the forehead	Wingtips stay near-black; the white forehead tuft never disappears; the old brass ring sits on the inside of the right wing; never swapped for an egret, an eagle or a phoenix

5 prop cards

Asset card	Requirement key	Script para graph	Structure and material	Lock clauses
Prop | Red-Wax Cross-Cloud Letter	prop_letter	P18	Aspect ratio about 1.5:1; the wax seal covers about a third of the flap; a hemp cord wraps once around it; 5 objects, 100 triangles	The seal is dark red, not bright red; the notch in the ridge pattern sits front and slightly right; no stamp, no writing
Prop | Brass-Shell Lamp	prop_lamp	P14 / P28	Hand-held brass lamp with a frosted glass chimney; wick colour 2200 K	The hand lamp in S1 and S5 and the tower backup lamp in S8 are built separately: similar silhouette, different construction
Prop | Skylark	prop_skiff	P14 / P24	12.2 m envelope; long shuttle-shaped wooden gondola; double envelope with three interwoven ropes and twin tail fins; 84 objects	Double envelope; the weaving order of the three ropes; tail fins in pairs; never a hot-air balloon or a propeller plane

Asset card	Requirement key	Script para graph	Structure and material	Lock clauses
Prop | Linen Mailbag	prop_mailbag	P22 / P30	Rough linen cross-body bag with a brass clasp and frayed edges at the mouth	Brass clasp; in S9 the burst clasp must be readable as one specific clasp; from S12 on the bag must be flat and empty
Prop | Tin Mailbox	prop_mailbox	P34	Old tin upright mailbox, the mouth worn bright, a thin steel plate at the bottom that gives the drop sound	Appears only in S11; the mouth faces the boardwalk; never a modern street mailbox

11 scene cards (parents) and their time-of-day cards

Scene parent card	Requirement key	Script para graph	State card (time of day)	Lock clauses
Scene | Cloud Ridge Post	scene_post_station	P14	Scene state | Cloud Ridge Post · early morning	The post always sits at the south-west end of the ridge and the lighthouse at the north-east end; the letter-drying line never moves; no modern buildings
Scene | Sorting Room	scene_sorting_room	P16	Scene state | Sorting Room · early morning	Window left, drawer right; a single oil lamp; bundles of old letters stacked in the foreground as a wall
Scene | East Ridge	scene_ridge_east	P20	Scene state | East Ridge · morning	The lighthouse stays in focus in the background; the ridge line matches the post; the direction the grass is pressed flat is fixed
Scene | Mooring Mast	scene_mooring	P22	Scene state | Mooring Mast · morning	Three mooring posts; a 14 m mast; the dock edge is the cliff itself, with no railing
Scene | Inside the Cloud Wall	scene_cloud_wall	P24	Scene state | Inside the Cloud Wall · noon	The cloud wall does not transmit light; the horizon stays tilted at 12° and is never corrected; rain lines are visible only in the searchlight

Scene parent card	Requirement key	Script paragraph	State card (time of day)	Lock clauses
Scene | Deep in the Cloud Wall	scene_cloud_wall_deep	P26	Scene state | Deep in the Cloud Wall · noon	Fog in three layers with the path closing again and again; no stable highlight except the crane's white forehead
Scene | Lighthouse Interior	scene_lighthouse	P28	Scene state | Lighthouse Interior · noon	Iron ladder left, lens centre, brass valve right; a faceted prism lens; wick at 2200 K
Scene | Inside the Skylark Gondola	scene_gondola	P30	Scene state | Inside the Skylark Gondola · afternoon	Three gondola windows; brass handrail on the right; crates may slide but never intersect
Scene | Cloud Rift	scene_cloud_rift	P32	Scene state | Cloud Rift · afternoon	The slit opens vertically, never sideways; the islet sits on the golden-section point; the light shaft lands only on the island and the sea
Scene | Islet	scene_islet	P34	Scene state | Islet · evening	Stone house, old mailbox and a wind-scoured boardwalk; a 14 m airship can moor at the end of the boardwalk
Scene | Above Cloud Ridge	scene_sky_ridge	P36	Scene state | Above Cloud Ridge · dawn	The post and the lighthouse can share the frame; the cloud deck is lit from a low angle; the post's window lamps come on one by one

Parent → state card reuse rule: All 11 state cards hang under their own parent card (parentKey points back to the parent), and **the parent card's primary reference**, changing only three things - time, colour temperature, cloud cover. Geometry, window and door positions, materials and furniture layout never move. This rule solves the most common failure: the same sorting room looking like two different rooms in S2 and S3. A shot binding points at "location + time of day", not at one generic scene image, so S2 and S3 bind the same "Sorting Room · early morning" and the audience sees one space.

Which cards each of the 12 shots binds

Bindings are not paperwork: every shot points its characters, props and scene at a specific card version, and a wrong pointer means the generation and finishing steps pull the wrong reference. The table is the real binding set, 53 entries in total; the count per shot is in brackets.

12-shot asset bindings (53 entries in total)

#	Character	Prop	Scene state card	Total
S1	Lan, Old Stationmaster	Linen mailbag, brass lamp, Skylark	Cloud Ridge Post · early morning	6
S2	Lan, Old Stationmaster	Red-wax letter	Sorting Room · early morning	4
S3	Lan	Red-wax letter	Sorting Room · early morning	3
S4	Lan, Old Stationmaster, Grey-Feather Cloud Crane	Brass lamp	East Ridge · morning	5
S5	Lan, Old Stationmaster	Skylark, linen mailbag	Mooring Mast · morning	5
S6	Lan	Skylark, linen mailbag	Inside the Cloud Wall · noon	4
S7	Lan, Grey-Feather Cloud Crane	Skylark	Deep in the Cloud Wall · noon	4
S8	Old Stationmaster	Brass lamp	Lighthouse Interior · noon	3
S9	Lan	Red-wax letter, linen mailbag, Skylark	Inside the Skylark Gondola · afternoon	5
S10	Lan	Red-wax letter, Skylark	Cloud Rift · afternoon	4
S11	Lan	Red-wax letter, tin mailbox, Skylark	Islet · evening	5
S12	Lan, Grey-Feather Cloud Crane	Skylark, linen mailbag	Above Cloud Ridge · dawn	5

Three things the binding table tells you:① S8 is the only shot in the film without Lan, and none of its three cards overlap with S7 - that is the physical evidence of the dual-line edit. ② The tin mailbox appears once, in S11, and a prop card used once is still worth building: it fixes the direction of the mouth and the metal thickness of the drop sound. ③ The most frequent cards are Lan (12 shots), the Skylark (7) and the red-wax letter (5), so those three get turntables and finishing first.

How each shot traces back to a script paragraph

The breakdown cuts the script into 38 paragraphs (P1-P38), and every shot points back at its own two: one titles the shot, one writes the action. This is not decoration; it is how you answer " who wrote the line I just deleted" when a shot changes. The left column is the shot, the right column the line that cannot be cut.

12 shots and their script paragraphs

#	Segment	Key line from the script
S1	P13-P14	"The Skylark rises from below the clouds and settles into the post, moisture on the envelope catching the morning light."
S2	P15-P16	"The old stationmaster takes the last cross-cloud letter from the back of the drawer and pushes it towards Lan."
S3	P17-P18	"The seal presses a ridge pattern into the wax, with a thin notch front and slightly right."
S4	P19-P20	"The grey-feather crane skims the ridge, wingtips almost brushing the stone."
S5	P21-P22	"The Skylark leaves the post. The gondola lifts first, then the wind pushes it towards the cloud wall."
S6	P23-P24	"The gondola heels inside the cloud wall, crates slide, the ropes tighten and hum."
S7	P25-P26	"The crane appears ahead of the bow, climbing into the wind, wingtips near-black."
S8	P27-P28	"The old stationmaster climbs down the iron ladder into the lamp room, opens the brass valve and strikes a match."
S9	P29-P30	"The clasp on the mailbag bursts and the red-wax letter slides onto the deck."
S10	P31-P32	"The cloud layer splits into a vertical seam and a shaft of light lands on a small island."
S11	P33-P34	"She pushes the letter into the slot and hears a very soft sound as it hits the bottom."

#	Segment	Key line from the script
S12	P35-P36	"Homeward. The Skylark climbs out of the rift above the ridge, the crane gliding beside the hull."

Eight stages, eight traces: every one opens in the platform

These eight stages are not a table of contents; they are eight real records inside the platform. Finish a film and all eight should hold something - if one is missing, the pipeline is broken. The table below is ordered "stage -> where -> what it leaves behind", and every number comes from the live project itself: breakdown confirmed at revision 2, an asset library of 30 cards, and a canvas of 111 nodes joined by 82 edges.

The eight stages, where they live and the measured counts

Stage	Where in the platform	What is left behind	How to check
1 Script	Episodic project " Letters from the Cloud Ridge" → script and breakdown	38 paragraphs (P1-P38), 25,219 characters; breakdown confirmed at revision 2, producing 30 requirements	Open the breakdown record, page to the latest revision and confirm status and timestamp in the footer; then open any requirement and check the script paragraph it cites
2 Shot list	Shot page of the same project	12 shots; 1752 frames of main film + 192 frames of titles = 1944 frames	Check frame count and duration per shot against this table; the totals must match
3 Assets	the asset-card chain table on the breakdown page + asset library + canvas asset-card nodes	all 30 requirements point at a card and a version; all 30 cards are locked, 3 / 5 / 11 / 11, each with one primary reference	expand the chain table on the breakdown page - all four counters should read 30; then filter the asset library by the four kinds and confirm every status is locked
4 3D previz	Canvas "Letters from the Cloud Ridge · 12-shot production canvas"	12 single-shot previz clips, 12 previz stills, 5 turntable sets totalling 480 frames	Nodes run in shot order; play each one and compare camera to the shot list
5 Layered rendering	Layered-comparison nodes on the same canvas	3 layered comparisons (S4 / S9 / S11) plus the six-collection switch demo	Open a comparison node and peel the layers: characters, props and set each on their own layer

Stage	Where in the platform	What is left behind	How to check
6 Final cut	Film nodes and platform export	720p master 81.004 s, platform 480p 81.0015 s, both 24 fps with AAC 96 kHz	Run ffprobe on both files, then play both all the way through
7 AI finish	Finishing nodes on the canvas	12 shot-level stills + 5 turntables + 1 poster = 18 images	Pair each with its previz still and check characters and props for drift
8 Course page	This chapter, the three language editions and the PDF	30-card contact sheet (1404×1263), shot list, binding table, acceptance table	Download the PDF and compare section by section; shot numbers and terms must match

Why insist that all eight are present: Script only and the film stays in a document; previz only and it stays a construction drawing; final cut only and the next character change starts from scratch. With all eight in place the film is re-workable: one card change traces to 12 shots, one script line traces to a specific shot. That is what this case most wants you to take away - **the final cut is only the ninth artefact; the first eight are why you can produce it again**.

Six collections and layered independent rendering

One project per shot; all objects are grouped into six collections by purpose: **01_Character / 02_Prop / 03_Set / 04_Light / 05_Camera / 06_Particles**. The point is not "to keep things tidy," but so that characters, props, and scenes can each be rendered independently and then composited back into the same frame. Items held or worn are still parented to the skeleton, but their collection assignment does not change-when the letter is in Lan's hand, it still belongs to the prop collection.

Collection boundaries and split rules

Collection	Contains	Never contains	Why render it separately
01_character	Skeletons and skinning for Lan and the old station master; the cloud crane's rig and feathers.	The letters, lanterns, and mailbags in their hands.	Swapping backgrounds, grading characters alone, building character mattes

Collection	Contains	Never contains	Why render it separately
02_prop	The red-wax letter, mailbag, lantern, cables, helm, and the entire Skylark.	Character bodies and scene rock volumes.	Letters and lanterns are brightened separately so they do not contaminate the environment; vehicles are output as separate layers.
03_set	The stone ridge, mail station, dock, sorting room, lighthouse interior, isolated island, and old mailbox.	Characters and carried props.	A clean background plate for atmosphere and depth of field
04_light	Morning light, window light, searchlight, lamp wick, lighthouse beam, and ambient sky light.	Any geometry	A light pass for checking direction and occlusion
05_camera	The single camera and its focal animation for this shot	A spare second camera	Makes camera and lens reproducible across 12 shots
06_particle effects	Cloud sea, fog banks, rain curtains, cloud-gap light shafts, and dust.	Character feathers and cloth.	Fog and rain are output as separate layers so density can be adjusted between the two storylines.

The six collections are **categories**, not "every shot must have all six." This production actually had two empty sets: **S6 has no characters and no set** (only vehicles, lights, and 7 effects objects), because this shot is a pure environment shot with "the ship but no people"; **S10 has no characters**, only the ship, the isolated island, and a single cloud-gap light shaft. This is not cutting corners-on the contrary, handing the weight to the vehicle and the environment leaves breathing room between the danger in S9 and the delivery in S11.

Collection composition and triangle counts for the 12 shots (measured).

Shot	Characters / Props / Sets / Lights / Camera / Effects	Triangles	Materials	Notes
S1	101 / 85 / 25 / 5 / 1 / 0	120,180	41	The heaviest shot in the film: two people + the crane + the whole station + the arriving airship.
S2	59 / 5 / 14 / 5 / 1 / 0	6,028	28	Indoor close-up, minimal geometry, carried by light placement.
S3	24 / 5 / 15 / 3 / 1 / 0	3,760	23	Macro shot, leaving only the hand, the letter, and the windowsill.

Shot	Characters / Props / Sets / Lights / Camera / Effects	Triangles	Materials	Notes
S4	72 / 0 / 25 / 5 / 1 / 0	39,736	28	No props; it all relies on the crane and the lighthouse.
S5	101 / 86 / 23 / 3 / 1 / 0	112,032	41	Leaving the station, with the airship and the dock in frame.
S6	0 / 85 / 0 / 3 / 1 / 7	58,048	18	No characters, no scene, pure vehicle + rain and fog
S7	72 / 85 / 1 / 3 / 1 / 5	46,520	24	The crane stands at the bow; fog is the protagonist
S8	24 / 0 / 9 / 3 / 1 / 0	3,456	16	The lightest geometry in the whole film; the narrative is completed by a single ignition
S9	35 / 6 / 20 / 6 / 1 / 0	9,776	23	Tilted interior, with the most lights
S10	0 / 85 / 15 / 3 / 1 / 1	83,372	23	No people, an isolated island and a light beam
S11	35 / 6 / 16 / 3 / 1 / 2	66,492	29	Delivery, a mailbox and a beam of light
S12	131 / 85 / 23 / 3 / 1 / 0	112,380	41	Return voyage: person + crane + boat + the whole station

What this set of numbers is for: Across all 12 shots the film totals 661,780 triangles. Seeing that S1, S5, and S12 are all above 110,000 and S8 is only 3,456 tells you which shots need rendering scheduled earlier and which can be re-rendered at any time **Look at geometry volume before scheduling**, far more reliable than guessing by feel that 'this shot seems complex.'



元素分层清单 · 每镜的对象归属

六个集合：01_角色 / 02_道具 / 03_场景 / 04_灯光 / 05_摄影机 / 06_粒子特效

镜头	段落	CHR	PROP	SET	LIGHT	FX	三角面	材质
S1	云脊邮站晨景	101	85	25	5	0	120,180	41
S2	分拣室	59	5	14	5	0	6,028	28
S3	蜡封与地址	24	5	15	3	0	3,760	23
S4	云鹤示警	72	0	25	5	0	39,736	28
S5	离站	101	86	23	3	0	112,032	41
S6	风暴云墙	0	85	0	3	7	58,048	18
S7	云鹤引路	72	85	1	3	5	46,520	24
S8	备用灯	24	0	9	3	0	3,456	16
S9	信件滑落	35	8	84	3	3	31,144	29
S10	云隙孤岛	0	85	15	3	1	83,372	23
S11	投递	35	6	16	3	2	66,492	29
S12	黎明返航	131	85	23	3	0	112,380	41

Six-collection element matrix. On the left are object categories; on the right are the collections and render passes they belong to. The purpose of this chart is to prevent ad hoc decisions about 'where this should go'-once an object changes collections between shots, layered rendering won't line up. Pay special attention to 'the letter': it is in the character's hand, in the mailbag, and in the mailbox, but across all 12 shots it always belongs to 02_Props.



Layered comparison for S9: character layer, props layer, scene layer, and composite result. Check three things-the contact point between the letter and the wooden floor, the occlusion relationship between Lan's hand and the envelope edge, and whether the direction of cabin-window light falling on the letter matches the light on the character. If any item is wrong, fix it in the project; don't force it in compositing.



Layer toggle demo (2.5 seconds): turn on character, props, and scene layer by layer. The acceptance criterion is 'turn off any layer and the remaining layers still hold up'-if turning off the scene makes Lan float in the air, the contact points weren't done right.



S11's layer toggles. This shot has two extra layers: light beam and grass blades. The light beam must be rendered separately so its angle and falloff can be controlled in post; the grass blades must be rendered separately so that when the beam sweeps past, they can hold the foreground without stealing focus from the face.

3D previz: solve space before texture

The preview stage answers only three questions: are the spatial relationships correct, can the camera position read the action, and does the lighting position separate the subject from the background? Materials, feathers, and fog details can all be added later; if the composition is wrong, it must be redone. This film's preview output uses the same camera positions, same focal lengths, and same frame counts as the final film, so the duration of every later shot doesn't need to change again. Render settings: 1280×720, 24 fps, real-time render engine, 16-24 samples per frame; 12 shots total 1752 frames, measured total render time **881.4 seconds** averaging 0.50 seconds per frame; the slowest shot, S8, is 0.75 seconds per frame, and the slowest single frame is 1.00 second. In other words **Re-rendering the entire final film once takes less than 15 minutes** -this is the precondition that makes 'preview first, polish later' viable.

12-shot render test (1280×720, real-time render engine, 16-24 samples)

Shot	Frames	Samples	Total time	Average per frame	p99
S1	168	24	80 s	0.48 s	0.67 s
S2	120	24	61 s	0.51 s	0.60 s
S3	144	24	61 s	0.43 s	0.62 s

Shot	Frames	Samples	Total time	Average per frame	p99
S4	144	24	104 s	0.73 s	0.92 s
S5	168	24	88 s	0.52 s	0.68 s
S6	144	24	88 s	0.61 s	0.75 s
S7	144	24	74 s	0.51 s	0.60 s
S8	144	24	109 s	0.75 s	1.00 s
S9	120	16	34 s	0.28 s	0.30 s
S10	144	24	48 s	0.33 s	0.37 s
S11	144	24	50 s	0.34 s	0.44 s
S12	168	24	85 s	0.51 s	0.95 s

When to stop and re-render: The preview stage can be rough, but these three things are not allowed-black frames, brightness jumps over 15% within a single shot, and action start points that don't connect to end points. Record all other issues in a notebook and fix them together during polish **Don't fix materials during preview**, because that is the most expensive rework: change a material once and all 12 shots must be re-rendered.

AI refinement: unify the texture, do not replace the shots

Preview solves space; polish solves texture. This film keeps the preview's camera positions, movement, and edit points, and does only four things to the image **Push the cool grade into a single cyan-gray hue, make the ochre-yellow cloak and dark-red wax seal the only two warm points, build layered aerial perspective for fog and clouds, and unify the film grain**. The approach is to produce one polished still per shot as a style anchor, then apply the same color parameters across the entire timeline; don't repaint shot by shot, to avoid the same Skylark turning into 12 ships across 12 shots.

[Project] 81-second Eastern sea-of-clouds postal-route short film, 16:9, realistic cinematic texture, cool cyan-gray base, slight film grain and thin fog layers, no text, no subtitles, no watermark, no logo.
[Character] Lan: around twenty years old, 1.7 m tall, short black hair pinned behind left ear with an old brass clip, ochre waxed-cloth cloak (hem faded white, rain-stained), dark teal high collar, leather wrist guards, coarse linen mail bag slung crossbody; do not change face shape, do not change height, cloak must not become bright yellow or orange-red.

[Character] Old Station Master: in his sixties, 1.58 m tall, slightly stooped, gray-white short beard, round copper-rimmed glasses, dark brown wool coat (shiny at elbows), carrying a copper-shell lantern; do not make him younger.
[Character] Gray-feathered cloud crane: standing height 2.2 m, wingspan 4.2 m, blue-gray feathers, near-black wingtips, a white tuft on forehead, an old copper ring on inner right wing; do not replace with egret, eagle, or phoenix.
[Vehicle] Skylark: envelope 12.2 m long, double-layer envelope bound by three ropes, long wooden spindle-shaped gondola, dark green paint with brass trim, paired twin vertical tails at rear; do not change to hot-air balloon or propeller plane.
[Prop] Red wax letter: beige envelope, aspect ratio about 1.5:1, dark red wax seal pressed at center of flap about one-third down, embossed with cloud-ridge pattern, notch on front slightly right, twine wrapped once horizontally; wax seal is not pure red, no stamp, no text.
[Scene] Cloud Ridge: a stone ridge 30 m wide and 64 m long floating above a sea of clouds, post station at southwest end (wooden main building, mooring mast, letter-drying lines), white stone lighthouse at northeast end; cool cyan slate, warm colors only on window light and copper fittings.
[This shot] S9 medium shot 35mm: gondola tilts sideways, letter on wooden floor slides along plank seams, Lan unbuckles safety belt and leans out, pressing envelope with fingertip an instant before it falls, shoulder hits cabin wall, cabin wall and mail bag as foreground.
[Movement] Camera fixed high inside cabin pointing down, tilts in sync with ship, envelope stays centered; no pan, no pedestal, no orbit.
[Consistency] Wax seal notch on same side; hair clip behind left ear; Lan's mail bag crossbody on right shoulder; in S8 the old station master lights the backup lamp inside the tower, not his lantern; no lighthouse light in cloud gaps in S10, only S11 turns it on.
[Negative] face swap, hairstyle change, cloak color change, wax seal turning pure red, single-layer envelope, single tail fin, camera shake, slow motion, extra characters, text appearing.

1. **Set one style anchor first:** Pick the shot that best represents the film's texture (in this film, S11: dusk, light beam, old mailbox), refine it until satisfied, then write that set of color, grain, and contrast settings as parameters and apply to other shots.

2. **Change one thing per shot:** In one shot, solve only one problem: either fog layering, warm-point position, or grain. Change three things at once and you won't know which one caused a problem.

3. **Re-check cuts after refining:** Refining changes brightness, so jumps at cut points must be rechecked. Keep the brightness difference within the two-line sequence (S7→S8→S9) under 15%; if it exceeds that, add a transition layer.

4. **Keep the previz version:** Save refinement as a new version instead of overwriting previz. When you rework a shot you need the spatially correct version back, not a rebuild from zero.



Refined poster frame: dawn sea of clouds, Skylark, and distant Cloud Ridge. Used for tutorial cover and film opening title, not part of the main edit. This image first establishes "how big the world is"; then the following 12 shots have a place to happen.



S1 refined still. Morning light presses in from side-back, sea of clouds below, stone ridge in middle, airship lifting from above; the three layers of atmospheric perspective must be countable. Warm color appears only at the window lamp, used to draw the eye to the two people in the porch.



S9 refined still. The envelope slides on the tilted cabin floor; Lan's hand is one finger-width away from it. The entire tension of this shot lies in "not quite there"-when refining, do not draw the finger and letter together; if that moment is caught, the shot is ruined.



S11 refined still. The letter drops into the tin mailbox, lighthouse light from right-behind passes through clouds onto her face. This is the only place in the film where light may fall directly on a face-if it happens earlier, the meaning of "being answered" is lost.



S12 refined still. The sea of clouds is smoothed by warm light, the post station lights are on one by one, Lan's mail bag is already empty. This is the only time warm color fills the frame, and the only "loop closure".



12 shot refined still contact sheet. Use it for the last thing: arrange the 12 thumbnails together, close your eyes and reopen them, see which shot "jumps out"-if one shot's color temperature is clearly different from its neighbors, go back and re-grade that shot, don't hide it in the edit.

Four layers of sound and two lines

Sound is built in four layers, from bottom up: **ambience** (high-altitude wind, rolling cloud sea, rain in the cloud wall, cabin chimney), **action** (cable strain, paper friction, the wax seal pressed by a fingertip, hollow footsteps on the gangway, one wingbeat, the helm's bronze shaft, that tin mailbox sound), **emotion** (Lan's two breaths, the old man's held breath when lighting the fire), **transitions** (the low hum of a sudden pressure drop, the "whoosh" of the wick catching).

This film has music-the opposite of Case 6. Case 6 relies on silence; this film uses a single motif to bind two lines together: **D minor plucked motif 4.5-14 s** (receiving the letter, restrained, slightly uneasy), **F minor brass 28.3-52 s** (from leaving the station to entering the cloud, kept in the low frequencies, never stealing the action), **S11 lighthouse high note lands at 66.6 and 67.9 s** (two hits, corresponding to "delivering the letter" and "the lamp lighting"), **dawn shifts to F major 70.2 s** until the end. The four nodes land exactly on the four narrative turns; where the music does not modulate, the emotion cannot change either.

> **How to decide whether to add music:** If the film can be understood by the cause and effect of action, music is a bonus; if the film relies on "music telling you to feel sad," then the music is doing the script's job. This film's approach is-the music appears only **between two actions**, never accents during the action. Try listening to that S9 section muted: wind, the rustle of the letter sliding across the plank, the dull thud of a shoulder against the bulkhead-none can be missing.

Platform film: one timeline, 14 video clips plus one audio track

The finished film does not use the automatic first cut; instead it is assembled by hand in the platform's final edit editor: **opening title card 0-4 s → S1 4-11 → S2 11-16 → S3 16-22 → S4 22-28 → S5 28-35 → S6 35-41 → S7 41-47 → S8 47-53 → S9 53-58 → S10 58-64 → S11 64-70 → S12 70-77 → closing title card 77-81**, plus one 0-81 s audio track. Video track 14 segments, audio track 1 segment, total length 81 seconds.

81-second delivery timeline

Segment	In	Out	Content	Duration
C0	00:00	00:04	Title card	4 s
S1	00:04	00:11	Cloud Ridge Mail Station morning scene	7 s
S2	00:11	00:16	sorting room handoff	5 s
S3	00:16	00:22	wax seal and cloud wall	6 s

Segment	In	Out	Content	Duration
S4	00:22	00:28	cloud crane warning, lamp out	6 s
S5	00:28	00:35	departure	7 s
S6	00:35	00:41	storm cloud wall	6 s
S7	00:41	00:47	cloud crane leads the way	6 s
S8	00:47	00:53	backup lamp (cut back to shore)	6 s
S9	00:53	00:58	letter slides away	5 s
S10	00:58	01:04	island in a cloud break	6 s
S11	01:04	01:10	delivery	6 s
S12	01:10	01:17	dawn return	7 s
C12	01:17	01:21	End card	4 s

The edit only does three things: sort by the shot list, cut on action, and breathe at the two black frames. One black frame after the opening title before the main film and one before the closing title card; everything else is hard cuts, with the cut point 3 frames before the action begins. The only exception is **S7 → S8**: this cut must land on the one frame of a crane wingbeat "down but not yet lifted"; the audience's breath is cut exactly there, and only cutting to the lighthouse gives the feeling of "changing places". S9 is not cut; let the list continue until the envelope is caught.

Acceptance: let the files speak

Delivery file measured parameters (ffprobe)

File	Resolution	Duration	Frame rate	Audio	Size
720p master	1280×720	81.004 s	24 fps	AAC, 96 kHz, stereo	53.5 MB
Platform 480p export	854×480	81.0015 s	24 fps	AAC, 96 kHz, stereo	10.0 MB

Completion criteria: The 12 shots play back to back; Lan's hair clip and cloak, the stationmaster's round brass glasses, the crane's wingtips and brass ring, the skiff's double envelope and paired tail fins, and the direction of the wax-seal notch all stay consistent across shots; the hairline crack in S3 and the letter dropped into the mailbox in S11 are the same letter; the lighthouse goes dark in S4, is lit in S8 and only reaches a face in S11, with no light at all in S10; Lan's mailbag is flat in S12; the breakdown is confirmed at revision 2, all 30 asset cards are locked with a primary reference each, and the 53 bindings point at card versions one by one; both the platform 480p and the local 720p master play from a direct link at 81.004 / 81.0015 seconds, 24 fps with a stereo track. Uploaded media alone, a single-shot candidate alone, or a queued export alone does not count as finished.

Do it yourself: rebuild this film in your workspace

1. **Create the blank canvas and fix its boundaries:** Create a blank canvas and write the boundaries down: title, 81 seconds, 16:9, 12 shots, no dialogue, two lines. Run the script breakdown first and let it list 3 characters, 5 props, 11 scenes and 11 scene states - 30 asset requirements; then fill in facts and lock clauses for each, upload one primary reference and lock it. The seven lead cards (Lan, Old Stationmaster, Grey-Feather Cloud Crane, Skylark, red-wax letter, Cloud Ridge Post, lighthouse interior) get turntables first; every other card needs at least one usable reference.

2. **Write the five acts and the 12-shot table:** Allocate time from the five-act table and the shot list above, then fill shot size, camera, lens, movement, key light and frame count. One main action per shot; looks go in the subject card, action goes in the shot description.

3. **First write the timeline of 'light':** Before writing the storyboard, write one sentence first: in which shot the lamp goes out, in which shot it is lit, and in which shot it falls on a person. This sentence determines the existence of S4, S8, and S11, and also determines whether the audience can connect cause and effect on their own.

4. **Build 3D previz:** Block the low-poly set, then fix camera, lens and movement before light positions. Render first, middle and last frames per shot, checking for black frames or jumps and keeping brightness variation within one shot under 15 percent.

5. **Split the file into six collections:** One project per shot; objects are grouped under 01_Characters / 02_Props / 03_Set / 04_Lighting / 05_Camera / 06_Particle Effects. Whether in a hand, in a bag, or in a mailbox, the letter always belongs to the props collection.

6. **Render four passes:** Render character, prop, set and a white-model pass separately, keeping light and camera visible in every layer. Composite three layers back on one frame and check edges, contact points and cast shadows.

7. **Pass the turnarounds first:** Lan, the Old Stationmaster, the crane, the Skylark and the red-wax letter each get a 96-frame turntable. A subject that fails its turntable does not enter final rendering. Even a prop card used by a single shot (the tin mailbox only appears in S11) must exist before you shoot it.

8. **Polish and unify texture, then do the dual-storyline edit:** Make one style anchor first, then change only one thing per shot; after polish, review the brightness at the three cut points S7→S8→S9. When editing, cut S8 in at the middle of S7's wing-flap.

9. **Lay in sound, export, and check item by item:** Lay the ambient bed first, then action sound, then music and the two lamp cues. Export 480p from the platform's manual edit, cut the 720p master from the same sequence, and use ffprobe to check resolution, duration, frame count, audio track and file size; then play both files through. Finally go back to the asset library and confirm all 30 cards are still locked and that the canvas node and binding counts have not dropped.

Asset and cost boundaries

□ Inside the platform: 30 locked asset cards (3 characters / 5 props / 11 scenes / 11 scene states) plus a 1404×1263 contact sheet, seven finished lead cards, 5 turntable sets totalling 480 frames, a 12-shot previz contact sheet, a six-collection element matrix, 3 layered comparisons with the layer-switch demo, 12 finished storyboard stills, 5 finished subjects and 1 poster image, 12 single-shot clips, the 81-second master and the platform 480p export - all published with this course for direct download and comparison. The brass lamp, linen mailbag and tin mailbox cards currently use full frames from S8, S9 and S11 for positioning; their single-object finishing renders come in a later revision - stated plainly here so you do not hunt for an image that does not exist yet.

□ Re-creation: Uploading, canvas organisation, storyboard structure, manual film, editing and export are available; generation models are billed per job, so check the quote before confirming and do not submit paid jobs that were not authorised.

□ Remix: You can reshoot your own version from the shot list; if you reuse this project's characters and sets, keep the attribution so readers do not mistake them for your own designs.

If you want to learn canvas operation first, go back to Case 2; if you want to practice single character and strong lighting, read Case 6; if you want systematic practice in shot language and dual-storyline editing, start from Advanced creation to begin.

Case 8: Full-flow rebuild • from script breakdown to "Snowridge Bells 2 • One Bell Returns"

What you see first: An 87-second, 12-shot snow film, "Snowridge Bells 2 · One Bell Returns", with a 1280×720 master, a 480p platform export, a 12-shot contact sheet and an 87-second timeline; the canvas also holds five turntables and 30 asset cards you can open for comparison. The 12 shots, the 87-second mix and the 720p master are the same files as Case 6; what this case adds is the 4,831-character script, a full breakdown into 30 asset needs, 30 locked cards, a 98-node canvas that carries all eight stages, and a fresh 480p export from the platform's Manual Film. So the thing to follow here is how needs grow out of the script, how cards get locked, and how each stage leaves a trace - not a different set of shots.

Studio → Episode project → Script → Script breakdown → 30 asset cards → 12 shot nodes → 3D preview → Layered render → AI finishing → Platform assembly → Export check

Live project: The film is a formal project in the studio: the workflow project "Snowridge Bells 2" holds the 87-second brief and the 4,831-character script (131 non-empty paragraphs). The script breakdown is confirmed at revision 2 and produced **30 asset needs** asset needs, mapped to **30 cards** cards, **30 locked V1 versions** locked V1 versions and **30 primary references** primary references. The four counts match with no drift; de-duplicated, those references are 19 distinct images, because several prop and state cards reuse a full frame from the parent card (see below). The canvas "Snowridge Bells 2 · 12-shot complete pipeline" has **98 nodes** nodes in eight stage zones. The film was assembled in the platform's Manual Film and exported: one 14-clip picture track plus one score track, checked after export. Open it after signing in: Snowridge Bells 2 · 12-shot full-flow canvas.

See the result: one master, one platform assembly



If playback fails, download the 720p master (1280×720, 87.000 s, 24 fps, AAC 48 kHz stereo, 39.2 MB).

The 87-second master of "Snowridge Bells 2 · One Bell Returns". Structure: a 4-second title card, 79 seconds of film (12 shots, 1,896 frames) and a 4-second end card, 2,088 frames in total. There is no dialogue and no melodic score: the sound is laid in four layers - ambience, action, emotion, transitions (the same four names case 6 uses for this film); the silence sits between the layers, it is not a fifth one.



If playback fails, download the 480p platform assembly (854×480, 87.000 s, 24 fps, AAC 48 kHz stereo, 4.46 MB).

This copy comes from the platform manual-film mode: the title card, 12 shots, end card and the 87-second mix are laid on one timeline, then Export preview -> Submit export, and the platform composites it. Both copies are 2,088 frames, and the 13 cuts found by frame differencing land on exactly the same frame numbers (96, 288, 432, 576, 720, 864, 1008, 1152, 1344, 1512, 1656, 1824, 1992); no shot was swapped. It is still not just a resolution drop: sampling mean grey every 3 seconds from 8-80 s (25 samples) gives 75.4 for the master and 97.3 for this copy, a ratio of 1.29, with the two curves still moving together (r=0.94) - the whole timeline was lifted. Whether that comes from the export pipeline colour handling or the source material, this tutorial does not attribute it; it only reports the measurement. Both cards are re-burned at their original pixel size: the ink bounding box of the title card is 541x67 pixels in both the master and the 480p copy; the end card is three lines with a 541x151 pixel bounding box. The 480p frame is smaller, so the same card takes up more of the picture and sits higher. Acceptance rests on ffprobe and per-frame sampling, not on file names, and not on the guess that only the bitrate should differ.



Twelve key stills in film order, one representative frame per shot. This sheet checks three things: whether S1 and S12 return to the same space (bookending, though the camera and lens differ - 35mm against 24mm), whether the three subjects' shadows in S8 fall on the same side, and whether the warm gold in S10 occupies only a small part of the frame (S12 is the one exception: there the warm orange floods the snow).



The 87-second timeline. The title and end cards are 4 seconds each; among the 12 shots S1 and S8 run 8 seconds, S9, S11 and S12 run 7 seconds, and the other 7 run 6 seconds, for 1,896 frames of main film. Note that S4 and S5 have the same duration but opposite screen direction - the only movement reversal in the film, used to show the bell rolling down while the fox drags it up.

One-line story and the seven-scene causal chain

In one line: **The old bronze bell on the snow pass gate has its rope worn through by wind and rolls downhill; a snow fox, a white camel and a condor take turns carrying it back up and hang it on a post on the rock shelf - when the bell rings, the storm stops.** The film has no dialogue, so causality has to live in the action: **Rope frays → bell falls → fox tests → fox drags → condor looms → three beasts shelter → split across the ridge → camel pulls → hung on the beam → three rings cross the valley → wind stops, sun rises**. Every link is a visible action; the audience is never told.

Seven-scene structure and causal chain (matching the 87-second film; timecodes summed from the 12-shot order)

Length	Timecode	Function	Key action	emotion
Title	00:00-00:04	Set the tone, give the title	Black frame; title and subtitle fade up (main title first, subtitle after), wind rises	Still
1 Stone gate	00:04-00:18	Establish bond and omen	The rope chafes on the lintel and snaps; the camel does not move, the fox looks at the rope first	Taut, about to break
2 Stone steps	00:18-00:36	First failure	The bell falls; the fox takes the rope and hauls it up, stopping halfway to breathe	Urgent, laboured
3 Rock wall	00:36-00:48	Blocked	The condor presses down with its shadow; wind drives the three beasts to the wall; the bell sits untouched between them	Pressure → pause
4 Snow ridge	00:48-01:03	Division of labour	The three stand one third each across the ridge; the camel bites the rope and pulls inch by inch	Silent, upward
5 Rock shelf	01:03-01:09	Payoff	The three animals haul the bell onto the crossbeam together and its heart lights warm gold for the first time; first strike at 01:06.5	One release
6 Snow valley	01:09-01:16	Result spreads	Second strike at 01:11.05; the three rings cross the stone gate, the snow ridge and the valley in turn; the wind drops	Open
7 Mountain pass	01:16-01:23	Leave an aftertaste	The cloud gap opens wider, warm orange pushes across the snow, the three animals become silhouettes; third strike at 01:17.4, leaving only low end and tail	Still
End card	01:23-01:27	Settling	Warm light levels the snow, the title fades out	Still

Why the warm colour is released in stages: Across S1 to S9 the snow sits entirely in cold blue-grey; in S10 the bell heart lights warm gold for the first time but still covers only a small patch (about 3% of pixels with R-B above 25); in S12, once the cloud gap opens, the warm orange finally spreads across the whole snowfield (about 8% by the same measure). The warm light grows a step at a time; it is not a single point for the whole film. One measurement worth recording: the three rings in S11 read as near-neutral and slightly cool in the finished cut (the brightest 1% of pixels average R-B about -11), and the warm gold stays at the bell in S10 - so do not assume a glowing bell ring is gold when you write prompts. Write that rule into every shot prompt rather than fixing it in the grade - a grade can unify colour, it cannot invent it.

The 12-shot list is the construction drawing

A shot list carries no adjectives, only parameters you can tick off: shot size, camera position and focal length, duration, frame count, action, sound design. The table below is the version actually used; the film totals 1,896 frames and 79 seconds and matches the finished cut frame for frame. With 96 frames of title and 96 of end card, the total is 2,088 frames and 87 seconds.

"Snowridge Bells 2" 12-shot build sheet (79 s of film)

#	Length	Shot size	Focal length	Duration	Frames	Action	Sound design
S1	Snow-ridge stone gate	Extreme wide to medium	35mm	8 s	192	Wind pours through the gate, snow grains blowing sideways. The camel stands outside, the rope taut on the old lintel beam; the fox crouches a step away, looking up at the rope.	Low wind throug h stone + grains striking rock + hooves pressing snow
S2	Snow-ridge stone gate	Close-up	85mm	6 s	144	The rope saws back and forth on the stone edge, fibres snapping one by one, the frayed end curling upward, the tail dropping into the snow.	Hemp fibres snapping one by one, 2 frames of silence at the break

#	Length	Shot size	Focal length	Duration	Frames	Action	Sound design
S3	Stone steps inside the gate	Medium	50mm	6 s	144	The fox sniffs the rope end, looks up at the ridge, raises a paw to point, then glances back at the camel.	Claws gripping ice + breath + distant wind pressure shifting
S4	Stone steps inside the gate	Wide trac king shot	35mm	6 s	144	The bell drops from the lintel, strikes the steps and rolls down the slope; the fox runs after the sound.	Metal impact + clapper rattling + dragging over snow
S5	Stone steps inside the gate	Wide	50mm	6 s	144	The fox takes the rope and drags it uphill, the bell carving a trench in the snow, wind flattening its fur.	Drag friction + panting breath + rising wind pressure
S6	Windbreak rock wall	Wide to medium	35mm	6 s	144	The condor comes down first as a darker shadow, then spreads its wings across half the sky; it makes no sound, only looks down.	Low airflow of opening wings, the bell falls silent
S7	Windbreak rock wall	Medium	50mm	6 s	144	First frame with all three: the camel kneeling, the fox sitting by its foreleg, the condor perched on the upper-left rock; their breathing slowly falls into one rhythm.	Wind recedes off-screen; only three breaths and falling snow
S8	Snow ridge	Extreme wide	24mm	8 s	192	Knife-edge snow line: camel left, fox right, condor on the middle rock, the bell half buried with only its mouth showing.	Steady high-altit ude wind, the bell occasionally lifted by gusts

#	Length	Shot size	Focal length	Duration	Frames	Action	Sound design
S9	Snow ridge	Wide	35mm	7 s	168	The camel lowers its head, bites the rope and drives all four hooves into the snow, hauling the bell up inch by inch until the rope is a straight line.	Rope fibres under load + hooves sinking + bell body dragging
S10	Rock shelf	Medium close	50mm	6 s	144	The camel brings the bell to the shelf, the fox braces its body with its head, the condor lifts the rope onto the beam with its beak; a ring of warm gold lights in the bell heart.	The rope drops onto the iron ring with a light knock; first strike at 01:06.46, then the bell heart lights
S11	Snow valley	Wide to extreme wide	24mm	7 s	168	The bell mouth throws out a ring of light, the halo forms, and three rings cross the gate, the ridge and the whole valley in turn.	Second strike (01:11.05): bell band +15.5 dB, low end +11.1 dB, 15 dB down 4.4 s after the peak
S12	Mountain pass wind gap	Extreme wide	24mm	7 s	168	The cloud gap opens wider, warm orange light pushes from behind the mountain, the snow turns gold layer by layer; the three beasts leave silhouettes on the shelf.	Wind stops; only the bell's tail and the faintest falling snow

Camera positions and movement live in the canvas 12 shot nodes and are not repeated here; the table only covers what you can compare at a glance. When something truly does not match, check the frame count first - matching duration with mismatched frames means the frame rate was changed.

Where each of the three strikes lands: The three strikes do not sit in the same shot: the first at 66.46 s (inside S10, 3.5 s after that shot starts), the second at 71.05 s (S11), the third at 77.42 s (S12), 4.59 s and 6.37 s apart. Measured in 20 ms windows over the 500-4000 Hz bell band, each peak sits 9.3 / 15.5 / 8.9 dB above its local floor; the low band (60-300 Hz) runs 10.7 / 11.1 / 20.0 dB above the bell band, so the third strike really is low end and tail only; 15 dB of decay takes 3.8 / 4.4 / 6.4 s. This row used to read "120 ms delay / 3.5 s mountain echo" for S11. That

was the intent in the design, not something measurable in the mix: raw-waveform correlation at +0.12 s and +3.5 s is -0.01 and -0.18, against 0.03 and 0.01 for the same test at 69.0 s. The column now carries measured values; if you want to keep a design figure, label it "design" instead of packing it into the same cell.

Script breakdown → 30 asset cards: the part of this case worth keeping

Case 6 worked the other way: it knew what it wanted and built seven subject cards. Case 8 takes a harder path: **Break the script into needs first, then let the needs generate cards**. Every noun the script writes - character, character state, prop, scene, scene state - becomes a need, and each need becomes a card; a card counts as done only when it carries a primary reference and a locked version.

This breakdown covered all 131 non-empty paragraphs and produced 30 needs, distributed as follows:

Distribution of the 30 needs (30 cards in total)

Category	Count	List
Character	3	White camel, snow fox, Andean condor
Character state	5	Loaded on the ridge, lintel hanging position, wary testing, rope in teeth dragging, looming silhouette
Prop	8	Old bronze bell, camel rope, saddle blanket, pack, lintel bell beam, wing-side copper ring, bell halo, old post on the shelf
Set	7	Snow pass gate, stone steps inside, windbreak rock wall, snow ridge, rock shelf, snow valley, pass wind gap
Scene state	7	Gate · before dawn, steps · storm, rock wall · night, ridge · morning, shelf · sunrise, valley · dawn, pass · wind stops

Four rules that are easy to get wrong:

1. **Needs grow out of the script, not out of your head.:** Each need carries quotes from the script plus paragraph numbers: the numbers count physical lines (170 in the full script, blank lines included), and the 30 needs here hold 57 quotes spanning P10-P147, every one of which matches the script verbatim. If a quote does not match, the breakdown ran on an older revision - re-read the script before running it again.

2. **State cards hang under a base identity; they are not a separate system.:** "White camel · loaded on the ridge" is a state of the white camel, not a second camel. The parent card owns what does not change (2.1 m at the shoulder, a notch in the left ear, mud on the neck); the child card only records what changes in this scene (weight over the forehooves, rope between the teeth).

3. **There is exactly one primary reference, and it must be verifiable in the image.:** Characters use turntables, props use turntables, scenes use clean plates, states use that shot key still. Without a primary reference the version cannot be locked - this is not process purity, it is that later finishing has no baseline to compare against.

4. **Locking does not mean "I looked at it", it means "it will not change".:** Once V1 is locked, every later shot uses it as the comparison baseline; to change it, cut V2 so the change stays traceable. Only when all 30 cards are locked and all 30 have a primary reference may previz start.

Acceptance looks at four numbers, and all four must match: **30 needs / 30 cards / 30 locked versions / 30 cards with a primary reference**, with zero drift. If any one differs, a card is missing, a version is unlocked, or a card has no reference. There is a stricter standard worth stating too: the 30 cards carry 30 primary references, but de-duplicated those are only 19 distinct images - the saddle blanket and the pack reuse the llama's full frame, the lintel beam reuses the stone gate, the llama-under-load state, the snow-ridge dawn state and the ridge share one image, and several state cards reuse the frame of their scene. The cards and the locks are complete, but "one dedicated reference per card" is not something this project achieved; do not go looking for 30 individual images.

Which cards each of the 12 shots should carry

Building the cards does not mean they are used. The table below is assembled shot by shot from the script: it lists which cards each shot should carry. That is a design intent, not a platform read-back. This project canvas only put each shot scene state into the storyboard node scene field; characters and props never got visual_binding edges (all 15 edges on this canvas are plain references; the per-shot binding used in Case 7 was not done here). Treat it as a checklist: every shot needs at least one card for what changes in that shot (a state card or a prop card), or the shot is just a scene display. To actually verify binding, draw the edges shot by shot on the canvas rather than reading this table.

Per-shot card table built from the script (design intent, not a platform read-back)

Shot	Scene / state	Cards for this shot
S1	Gate · before dawn	White camel / snow fox / camel rope / lintel bell beam / snow pass gate

Shot	Scene / state	Cards for this shot
S2	Gate · before dawn	Camel rope / lintel bell beam / snow pass gate
S3	Steps · storm	Snow fox · wary testing / snow fox / white camel / camel rope / snow ridge / stone steps inside
S4	Steps · storm	Old bronze bell / lintel bell beam / snow fox / stone steps inside
S5	Steps · storm	Snow fox · rope in mouth, dragging the bell / snow fox / old bronze bell / camel rope / snow ridge / stone steps inside
S6	Rock wall · night	Condor · looming silhouette / Andean condor / lintel bell beam / snow fox / windbreak rock wall
S7	Rock wall · night	White camel / snow fox / Andean condor / old bronze bell / windbreak rock wall
S8	Ridge · morning	White camel · loaded on the ridge / white camel / snow fox / Andean condor / old bronze bell / snow ridge
S9	Ridge · morning	White camel · loaded on the ridge / white camel / snow fox / old bronze bell / camel rope / snow ridge
S10	Shelf · sunrise	White camel · lintel hanging position / white camel / snow fox / Andean condor / lintel bell beam / rock shelf
S11	Valley · dawn	Bell halo / snow pass gate / snow ridge / snow valley
S12	Pass · wind stops	White camel / snow fox / Andean condor / old bronze bell / lintel bell beam / rock shelf

Two rules for reading the table: the saddle blanket and packs ride on the camel from S1 onward (P30-P35 say the packs shift its weight right), so they are not repeated for every shot and appear only in the S8 White camel loaded on the ridge state card; the condor **wing-side copper ring** is part of the character own identity (P18), not a prop for one particular shot, so it travels with the character card.

Eight stages, eight traces: every one opens in the platform

The eight stages are not eight folders; they are eight places that must leave a trace in the platform. Miss one and the film becomes single-use: with only a script it stays a document; with only a finished cut, changing one character means starting over.

Where the eight stages live (actual counts for this project)

Stage	Platform location	What this project left there
① Script	Episode → Script	4,831 characters / 170 lines / 131 non-empty paragraphs, status draft, revision 2
② Shot list	Canvas zone ②: full table + 12 shot nodes + 12 key stills	12 shots / 1,896 frames / 79 s, cut points and sound design aligned shot by shot
③ Assets	Episode → Script breakdown and asset cards; canvas zone ③ with 30 cards	30 needs → 30 cards → 30 locked V1 versions → 30 primary references
④ 3D preview	Canvas zone ④: 12 single-shot clips + 5 turntables	12 clips totalling 79 s, camera, focal length and duration matching the film
⑤ Layered render	Canvas zone ⑤: S4 / S11 comparisons + a 52-second layering demo	Six collections × four passes, the bell on its own layer
⑥ Film	Canvas zone ⑥ + the platform Manual film timeline	720p master, 480p platform export, 87-second mix, title and end cards, poster
⑦ AI finishing	Canvas zone ⑦: one per-shot prompt set covering all 12 shots + 5 refined frames	The prompt is one 5,510-character text split into 12 sections from 【S1 Stone Gate and White Camel】 to 【S12 Dawn Resolution】, sharing one set of consistency anchors and negative prompts; the refined frames are S1 / S3 / S8 / S9 / S12
⑧ Course page	Canvas zone ⑧: eight steps + acceptance checklist	Parameters, pitfalls and acceptance criteria, all reproducible

Layered rendering: six collections, four passes

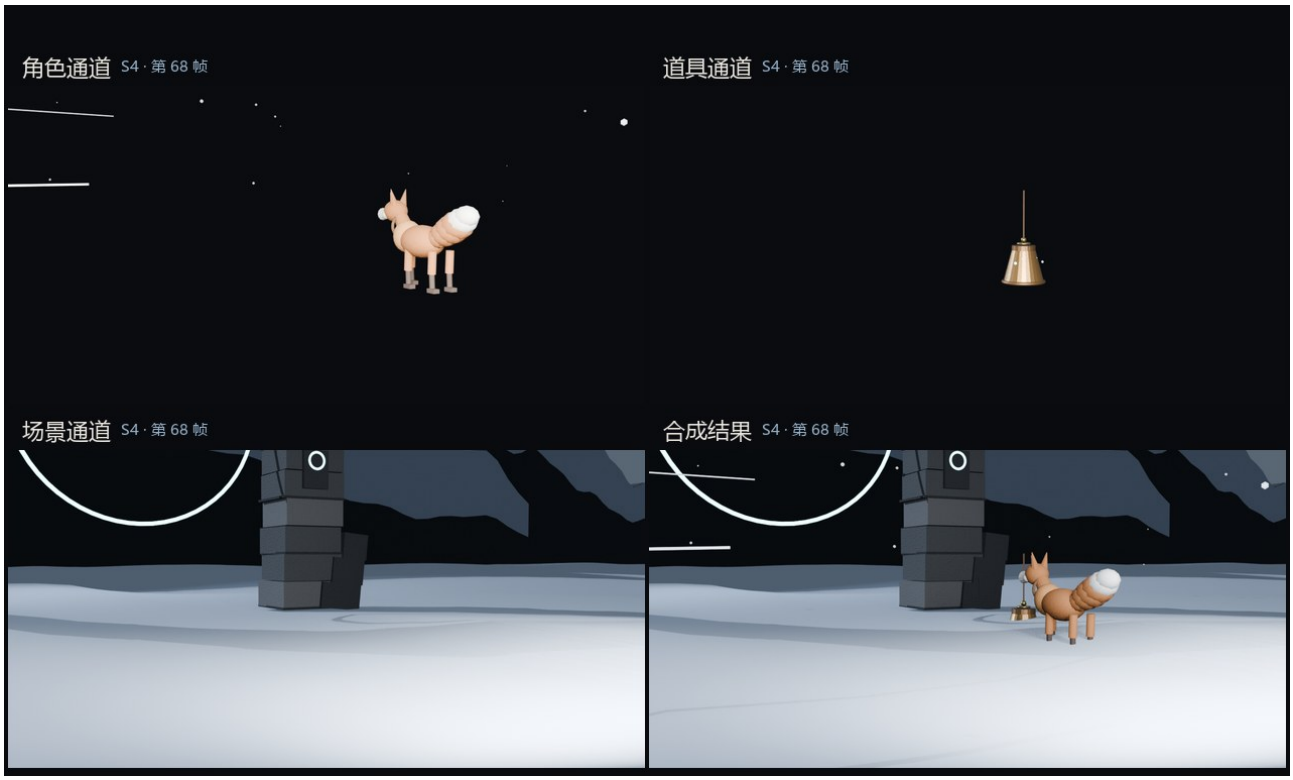
Layering is not for looks; it is so you can redo only the layer that broke. Each of the 12 shots has its own project, with objects named into six collections: **01 characters (white camel, snow fox, condor) / 02 props (bell, rope, saddle blanket, pack) / 03 scene (gate, steps, rock wall, ridge, shelf, valley) / 04 lights / 05 camera / 06 particles (snow grains, wind streaks, bell halo, mist)**. With one camera and one lighting setup, toggling collection visibility alone yields four passes:

□ Character pass: characters only, contact shadows kept;

□ Prop pass: the bell on its own layer, with glow strength on a separate control;

□ Scene pass: no characters, contact points and shadow positions kept;

□ Clay pass: low-saturation materials, used to confirm composition and silhouette.



The S4 comparison. The thing most often missed here is the bell shadow - as the bell rolls over snow the shadow has to change with its rotation, not stay a fixed ellipse. So props go out on their own layer and the shadow gets a layer too.



The S11 comparison. The three rings are three independent layers, evenly spaced and offset in time; when the glow stacks it must not eat the valley shadows, or AI finishing will generate noise there.

When stacking back, check three things per shot: **Edges** (fur, feather tips, frayed rope), **Contact points** (hoof soles, claws, rope meeting the beam), **Shadows** (the three subjects shadows pointing the same way). If any one fails, re-render that collection alone instead of the whole shot.

3D previz: solve space before texture

Every shot gets one animated preview with camera, focal length and duration aligned to the film frame for frame; five turntables (camel, fox, condor, bell, gate) fix scale and structure and never enter the film. Preview answers only three questions: are the start and end frames right, do subject size and shot size match, and is the direction of motion continuous with the neighbouring shot.

The easiest thing to get wrong is direction: in S4 the bell rolls down, in S5 the fox drags it uphill, and the camera stays below the slope, so on screen they move opposite ways. Shoot S5 downward too and the audience reads the fox as pushing the bell down. S8 is the film only rule-of-thirds frame - three subjects a third each, shadows on one side, about twice body length.

AI refinement: unify the texture, do not replace the shots

The goal of finishing is to make 12 shots look like one day on one camera, not to make each shot prettier. So every prompt carries the same consistency anchors; when you change a shot you change only the shot paragraph and not one word of the anchors. Change one variable at a time: this pass only snow density, the next only fur lift - never composition and light direction together.

[Project] 87-second snow film, 16:9, realistic cinematic look, low-saturation cold blue-grey, light film grain, no text, no subtitles, no watermark, no logo.
[Consistency anchors] White camel: 2.1 m at the shoulder, creamy white coat, dark grey-brown slush on neck and knees, a notch in the left ear; snow fox: 0.75 m body, grey-white winter coat, near-black tail tip, right shoulder half an inch lower than the left; Andean condor: 1.2 m standing, 3.0 m wingspan, grey-white ruff, a ring of old copper on the inner right wing; bronze bell: 18 cm tall, bell-shaped, a chip in the mouth, bone clapper.
[Unified lighting] Snow key is cold blue-grey (about 6800K); warm colour has only two sources: the bell heart (about 2400K, from S10 on, a small patch) and the warm orange after the cloud gap in S12; shadow direction is consistent throughout.
[This shot] {shot} | {shot size} | {focal length} | {duration} | {action}
[Camera and movement] {camera and movement}
[Material notes] Snow grains have volume and visible texture; fur lifts rather than slides; bronze has oxidation and hand-polished highlights; hemp rope is frayed.

[Negative] no text, no watermark, no logo, no HDR blowout, no plastic fur, no smooth CG skin, no extra characters, no modern objects, no warm spill, no lens distortion.



S9 retouched frame. Again exactly three checks: the bell is not yet lit (warm gold only arrives in S10), every subject's shadow falls to the right of frame, and the rope drag mark on the snow is continuous. Only when all three pass may this version go back into the sequence.

Platform assembly: one timeline, 14 picture clips + 1 sound clip

The film was laid out by hand in the platform, not composited locally and uploaded. Two tracks: **Picture** Picture 14 clips (title 0-4, S1-S12 running 4-83, end card 83-87), **Music** Music 1 clip (the 87-second mix across 0-87; ambience, action and bell only, no melody). Check total duration and clip count in Export preview, then submit.

The platform Manual film timeline

Track	Clips	Content	Timecode
Picture	14	Title card 1 + 12 shots + end card 1	00:00-01:27
Music	1	87-second mix (ambience, action, bell; no melody), across the whole film	00:00-01:27

Manual film is not a licence to guess for free: Manual assembly and export cost no generation points, but you still go preview → submit, and each submit carries the change token returned by that preview; if the token no longer matches the project you get 409, and the fix is to re-read the project rather than retry the same request. Actual spend for this export was 0 points.

Acceptance: let the files speak

Acceptance record for this case

Item	Measured	How it was measured
Master	1280×720, 87.000 s, 2,088 frames, 24 fps, H.264 + AAC 48 kHz stereo, 39,199,564 bytes (39.2 MB)	ffprobe for container and streams, plus -count_frames for the frame count; the video stream md5 is identical before and after the fix, only the audio was replaced
Platform export	854×480, 87.000 s, 2,088 frames, 24 fps, H.264 + AAC 48 kHz stereo, 4,458,231 bytes (4.46 MB)	ffprobe and -count_frames again on the exported file
Timeline	2 tracks / 15 clips / 87.0 s total	Clip count and duration returned by export preview
Breakdown chain	30 needs / 30 cards / 30 locked versions / 30 primary references, drift 0	Read back the breakdown revision + count the asset library
Canvas	98 nodes across eight zones; 15 links (script → bios → beats → shot list → 12 shots)	Read back the canvas document
Assets	62 items: 61 uploaded media assets + 1 platform-composited film	Asset list production method column
Audio continuity	Master: 4,080 audio packets, largest gap between neighbouring packets 0.021 s (one AAC frame), no holes at all; the platform export also has 4,080 packets and no gap, and both line up sample for sample with the 87-second mix	ffprobe reading packet=pts_time,duration_time to find gaps in the time difference between neighbouring packets, then decoding both audio tracks to 48 kHz float PCM and subtracting them sample by sample

Three things you must look at: is the title card 4 seconds, is there a 4-second end card after S12, and does the film have an audio track. Then add two machine checks: **does the audio track have holes**, and **does the sound sit out of sync as a whole**. The total duration of 87.000 s and an audio stream of 87.000 s line up, but that does not mean every segment carries sound, or that the sound still sits on the picture. The master shipped here first hit both

problems at once: not a single audio packet between 3.99 and 6.92 s, about 2.93 s of digital silence; and after that hole the whole track ran 2.93 s early, so the third strike that belongs in S12 (01:16-01:23) arrived at 01:14.5, inside S11. The fix needs no re-render: remux the 87-second mix into the master without re-encoding video - the video stream md5 is 95243c83..., both before and after, and only the audio changes from 3,943 packets to 4,080 packets that match the mix sample for sample. After the fix the three strikes sit at 01:06.46 / 01:11.05 / 01:17.42, matching the shot list one for one. **Put each strike next to the shot list** is the cheapest check here: when a timecode does not match, the sound moved, not the picture. Numbers that add up while the title card has stretched to 5 seconds is the most common kind of "looks fine"; timestamps that run continuously with a 2.9-second hole in the middle, or a track shifted as a whole, is the second kind.

Completion criteria: Done means you can walk from script revision 2 through to 30 locked asset cards, with four numbers (requirements / cards / locked versions / primary references) all equal to 30; every shot duration and frame count line up on the canvas, all 30 primary references open from the asset library, and 1,896 frames of main film plus the two cards make 2,088; each of the six collections renders on its own, and the edge, contact and shadow checks pass on the character, prop, set and clay passes; the platform timeline holds 2 tracks and 15 clips over 87.0 seconds, the 480p export is 854x480 at 24 fps with a 48 kHz stereo track, and its 13 cuts land on the same frame numbers as the master; warm light grows from small to large - only a patch of gold in the bell heart in S10, the whole snow face only in S12; and all three bell strikes land after the bell is hung at 01:03. Cards finished, a single-shot candidate, or an export still queued all fail.

Do it yourself: rebuild this film in your workspace

1. **Write the script until you can count paragraphs.:** 4,000-5,000 characters and about 131 paragraphs works best: too short and it will not yield 30 needs, too long and the shot list runs away. When it is done, count the named things yourself.

2. **Run the breakdown and trust only the quotes.:** Every need must click back to its paragraph. If a quote does not match, do not rush to edit cards - first confirm the breakdown ran on the current script.

3. **Put a primary reference on all 30 cards and lock them.:** The four numbers must be equal. Skip this and every shot drifts later.

4. **Build the 12-shot preview and nail the durations.:** 1,896 frames in total. Changing duration is cheapest at preview; at the film stage it means re-laying the whole audio track.

5. **Layer by the six collections and re-render only what broke.:** Three checks: edges, contact points, shadows.

6. **Write the 12 per-shot prompts in one long text and leave the anchors alone.:** Split into 12 【S# shot name】 sections sharing one set of consistency anchors and negative prompts; change one variable at a time, then compare side by side against the primary reference.

7. **Assemble in the platform, preview first, then submit.:** 87 seconds total, 14 picture clips + 1 sound clip. On 409, re-read the project instead of retrying.

8. **Run ffprobe immediately after export.:** Duration, resolution, frame rate and audio track - missing one means not done.

To learn canvas operations first, go back to Case 2; to see the same shots as a finished film, read Case 6; to practise two-strand storytelling, read Case 7; for systematic work on shot language, start with Advanced creation.

Asset and cost boundaries

Every frame in this case was made inside this project: script, breakdown, cards, preview, layering, finishing and film all leave traces in the same space, with no outside footage. 3D preview and layered renders feed the film directly; AI finishing only unifies texture and never replaces character or prop structure. Manual assembly and export in this case incur no generation cost; actual spend was 0 points.

05 Advanced creation
story, shots, continuity, colour, edit and review

Advanced creation: Make better videos

Once you've made your first film, learn to answer three questions first: Why shoot it this way? What's not working? How do you fix it with the fewest reworks? The six chapters here cover creative judgment; the feature manual covers buttons and operations. Use them together.

Using "The Last Lamp," prepare an existing short film, shots, or your own materials and practice on a copy. If you have no materials, complete a written storyboard and review checklist first; you don't have to pay to generate again just to learn each item. Work through the first five chapters in order, then use Chapter 6 to complete one round of targeted revisions; you can also look things up by problem.

> The good/bad comparisons below are original written storyboards and edit plans, not comparative tests of real generated footage. Rules are tools that help an audience understand, not formulas you may not break; when you break one deliberately, say what you want the audience to experience. Model output still has to be watched; a prompt, a reference or a successful job does not guarantee artistic quality.

How to study each chapter

1. Answer the core question, understand the principle, then look at the comparison for the same story.

2. Go back to TimeGleam and make one small change, keep the original version, and don't make "more generation" the goal of the exercise.

3. Use the acceptance checklist to state what improved; if it failed, go back to the problem instead of piling on more style words.

> Deliverables for this piece: a shot intent sheet, a continuity log, a two-version pacing comparison, and a prioritized revision list. Advanced exercises do not count toward progress in the original eight lessons and do not automatically submit generation jobs.

Advanced 1: Story and emotion

Core question: Why should the audience keep watching, and what has changed by the end compared with the beginning?

Understand the principle first

First write in one sentence: "what the character wants-what stops them-what choice they make." The obstacle doesn't have to be a villain; it can be hesitation, time, or a lack of information. A short film doesn't need to cram in a full feature structure, but it should at least make the audience feel a change. Emotion comes from visible action and trade-offs, not from writing "moving" or "shocking" in the prompt.

In The Last Lamp, you can use "hurrying to shelter from the rain" and "deciding to take the keychain home" as the before and after states; the word "home" is new information, and the pause before touching it is the process of choosing. If you only show a pretty rainy street, character, and prop, there is no causality between the three; viewers see only a set of style frames.

Good vs. bad: text storyboard

Compare information and action, not prompt length

Likely to fail	Clearer approach
"He misses his hometown deeply," but the frame keeps showing the character walking.	First discover the keychain, then make out "home," then pause and grip it; the action responds to the new information.
Every shot captures rainy-night atmosphere, then the prop is suddenly removed at the end.	Keep one environmental shot to establish the rainy night, and give the runtime to discovery, recognition, and decision.

TimeGleam hands-on practice and exercises

Text node → save story intent; Shot draft → confirm shots; Timeline → compare against the preview

1. In the Text node, write the opening state, trigger information, and ending choice clearly. Keep 1 character, 1 location, and 1 prop; do not let the assistant add characters to explain things on its own.

2. For each of the existing 5-6 shots, write one sentence: "What does this shot let the audience know/feel?"; delete or merge repeated information; revise the text and order first, without generating.

3. Play the existing clips muted and check whether "discovery-hesitation-choice" reads clearly. If the engraving is illegible, first consider a clear prop close-up or proper subtitles; do not rely on blurry model text to carry key plot.

Acceptance checklist

□ Without explaining the backstory, viewers can still identify the character's starting point, the trigger information, and the choice.

□ Each shot adds new information or necessary emotion; it is not merely another good-looking repeat.

□ If turning off the lights makes viewers misread the character's identity, you can state whether to keep that implication or change it to leaving the store directly.

Review: Ask someone to retell what happened after watching, record the gap between their understanding and your intent, and do not explain the film for them first.

Next chapter: Turn story intent into shot choices.

Advanced 2: Storyboards and visual language

Core question: Where should this shot make the audience look, and why does it need this shot size, camera position, and movement?

Understand the principle first

Storyboarding is not about dividing the text into six equal parts; it is about arranging the order in which information appears. Wide shots are good for establishing the relationship between a character and the environment, medium shots make action easier to observe, and close-ups can emphasize expressions or prop details; they are not fixed emotion buttons. Camera height and distance should serve the current perspective, not every shot needs a high angle, low angle, and "cinematic feel."

First decide the center of attention, then choose background layers and negative space. Camera movement must have a reason: follow an action, reveal new information, or change the relationship between the character and the environment. A static shot can also have tension; when the character's action is already complex, adding multiple camera moves increases the difficulty of expression and generation.

Good vs. bad: text storyboard

In the same "discovery" beat, different shots do different jobs

Likely to fail	Clearer approach
Using only wide shots keeps the key prop unclear; then a sudden close-up gives no sense of where it is.	Use a medium shot to establish the character's relationship to the counter, then cut to a prop close-up, and return to a character close-up for the reaction.
One shot pushes in, orbits, and pulls out while the character picks it up, examines it, and leaves.	This shot only slowly pushes in to the discovery; examining the prop becomes its own shot, and leaving is saved for the ending.

TimeGleam hands-on practice and exercises

Shot draft → shot size / action / camera movement; Image or video node → prompt and actual reference

1. Set up three lines of shot intent: medium shot "establish the character and counter," prop close-up "read the message," character close-up "see the hesitation." Write one main action per line.

2. Write two options for S04: A locked-off camera, B slow push-in; change only the camera move and keep the character, light and references identical. Compare the written intent first, and compare playback only when footage already exists.

3. Spatial previs can be reviewed, Panorama and 3D director's desk and you don't need to make 3D or storyboards for every shot. For actual generation, submit only the approved option; don't generate two sets by default just for comparison.

Acceptance checklist

□ Pause any shot, point out the primary focus of attention, and explain whether the background steals focus.

□ Key actions and objects are legible at the expected viewing size, not only when zoomed in.

□ You can explain the choice of locked-off or moving camera, and don't add camera movement just for show.

Review: With the prompt covered, looking only at the frame, does the actual attention match the intent?

Next chapter: Make adjacent shots belong to the same space and action.

Advanced 3: Continuity and consistency

Core question: After a cut, does the audience still know it's the same person, the same place, and the same action?

Understand the principle first

Appearance consistency is only the first layer; you also need to track wardrobe, prop state, which hand is holding what, position, direction of movement, and timing. A continuity record captures "what has already happened," for example whether the keychain is still on the counter or already in the right hand; you can't just store the character master image.

The axis can be understood as an imaginary line between a character and the object they're observing. Continuous editing usually places the camera on the same side of the line to reduce sudden left-right flips; eyeline matching also requires checking the height and target position of the gaze. For action continuity, check whether the end of the previous shot and the start of the next are at compatible stages, so the same reach isn't repeated twice. When you need to cross the line, first use a shot or visible movement that re-establishes the space so the audience can follow; don't treat the rule as an absolute prohibition.

Concept reference: Tom Barrance: diagram of the axis, eyeline and space (learnaboutfilm.com/film-language/sequence/180-degree-rule/). Story examples and exercises are original.

Good vs. bad: text storyboard

Check the cut first, then whether the individual shot looks good

Likely to fail	Clearer approach
The character looks down-right in the previous shot, but the prop's position in the next shot implies it's up-left.	First establish the counter and prop positions in a medium shot, then shoot the reaction and close-up using the same spatial relationship.
In the previous shot the right hand has just grasped it; in the next shot the prop is back on the counter and the left hand reaches for it again.	Record the holding hand and action stage, choose compatible cut points, and add a neutral prop close-up if needed.

TimeGleam hands-on practice and exercises

Subject/actual reference → shot draft → assets and versions → current timeline

1. Use text nodes to record the character's appearance, the positions of the counter and door, prop positions, which hand is holding what, and the light direction. The subject image handles appearance; the text record handles state. The two can't replace each other.

2. Check the end and start of adjacent shots and record any inconsistent positions or actions. References must actually be added to the node; writing only "same as previous shot" won't automatically send the previous shot to the model.

3. First try trimming, switching to an existing candidate, or adding a valid insert shot; only redo the specific shot if it still doesn't work. Don't blindly mirror props with text just to fix direction, or the text and handedness will flip.

Acceptance checklist

□ Characters and props not only look the same; their positions, holding hands, and states also match across shots.

□ Movement direction, eyelines and spatial relationships between adjacent shots can be explained; intentional jumps have a narrative reason.

□ After swapping a node candidate, the timeline's actual references have been rechecked; a new canvas result is not treated as an updated edit.

Review: Play only the two junction points - does it read as one continuous action?

Next chapter: Keep light and color on one visual logic.

Advanced 4: Light, color and art direction

Core question: Why is the light and color in this frame arranged this way, and does it help the audience understand the character?

Understand the principle first

Establish a reason for the light source first: the counter lamp lights the hands and the side of the face, while the rainy night outside the window provides ambient color. Then look at light and dark, warm and cool, and foreground-to-background layering so the subject separates from the background. Soft and hard light describe the transition between light and dark and the texture of shadows; they do not mean "good" and "bad" - the choice should serve the scene and the emotion.

Manage the whole film with a few stable conventions: primary and secondary colors, costume materials, prop wear, light direction, background complexity. The warm shop lamp and the cold rainy night are design choices for this film, not the general conclusion that "warm colors are always happy, cool colors always sad." Local variation can change when there is a reason; consistency does not mean bringing every frame to the same brightness.

Good vs. bad: text storyboard

Unify the visual logic, not apply the same filter to every shot

Likely to fail	Clearer approach
Adding neon, cyberpunk, film, dreamlike and other competing keywords to every shot.	Fix "old wooden counter, warm lamp, deep blue rainy night, restrained saturation," and add only the changes the current shot needs.
The lamp is on the character's left in the earlier shot, but with no spatial change it becomes hard light from the right in the next.	Write the positions of the lamp and the window into the visual conventions, and check shadows, reflections and background relationships shot by shot.

TimeGleam hands-on practice and exercises

Text / subject reference kit → real references; canvas settings / node settings → local style; timeline → effects and color grading

1. Write six lines of visual conventions: primary and secondary colors, light source, costume, counter material, prop state, modern decor that must not appear. Check them against references instead of stacking keywords.

2. Pick three existing stills and compare them side by side, noting differences in subject brightness, light direction, color and props. There is no need to generate three new images for the exercise.

3. Try local adjustments in the edit first for slight exposure and color differences; when the light structure, character or props are wrong, filters and upscaling cannot replace a redo. Experiment only on a copy of the canvas or node so the shared defaults for the whole work are not changed by mistake.

Acceptance checklist

□ Three shots side by side read as the same scene, and the character and key props are easy to identify.

□ You can point to where the light comes from and why the colors were chosen, not just answer "cinematic."

□ Preview on the intended device: faces are not darkened to make the warm light stand out, and lamp and window detail is not wiped out to brighten the image.

Review: Ignore detail sharpness for now - are the light and color of this shot already helping it communicate?

Next chapter: Shape time and emotion with editing and sound.

Advanced 5: Editing rhythm and sound

Core question: When do you cut, which sound arrives first, and how long does the audience need to stay?

Understand the principle first

Editing is about allocating attention and time, not dividing shot length evenly. Establishing the environment can be brief, the audience needs time to read key prop information, and the pause before a decision can be kept. Decide length by whether the information was read, the action is complete and the emotion has happened; do not treat "5 seconds per shot" or "cut on every beat" as a formula.

Separate the ambience, action sounds, and dialogue the characters can hear from the added score the audience hears. Rain can connect spaces, a doorbell and touch sounds can emphasize action, and the score should carry emotion with intent. Make key sounds clear first, then place background, and don't let music bury story information.

Sound and picture don't have to cut together: a J-cut brings the next shot's sound in before its picture; an L-cut carries the previous shot's sound into the next picture. Try letting the doorbell lead into the shop entrance, or letting the tail of the shop ambience run into the character's reaction; if it feels unnatural, undo it-don't use a technique just to use the term.

Concept reference: Adobe: J and L cut editing (helpx.adobe.com/premiere/desktop/edit-projects/trim-clips/perform-j-cuts-and-l-cuts.html) borrow the principle only; don't treat Adobe buttons as TimeGleam features.

Good vs. bad: written editing plan

Establish rhythm first, then consider decorative transitions

Likely to fail	Clearer approach
Keeping dead waiting time to hit 30 seconds, then cutting too fast when reading the engraving.	Cut the waiting and give the time to information and decisions; let total length serve the expression instead of forcing a set number of seconds.
Every shot's location sound cuts off abruptly, and every cut gets a flashy transition and big music.	Try straight cuts first, handle ambience continuity and any necessary fades, then decide on transitions.

TimeGleam hands-on practice and exercises

Current timeline → in point/trim/audio track/volume and fades → preview

1. Keep the original project and work on a copy: make version A with even shot lengths, and version B with durations allocated to "discover-recognize-hesitate-decide." Reuse existing assets and change only trims and arrangement.

2. Use standalone rain, doorbell, or location-sound assets you have the rights to and try one split edit, placing the sound and picture clips separately. If you don't have standalone sound, write a sound list first; don't assume the current interface can automatically separate any video audio track, and don't call a music model by default.

3. Watch the story muted first, then listen to sound only, and finally play them together. Adjust in points, volume, and fades with the actual controls; don't use uniform speed changes to patch every action and lip-sync problem.

Acceptance checklist

□ Every hold in version B has a reason, and you can name at least two cutting decisions that differ from version A.

□ No unintentional repeated reaching, abrupt sound cuts, or key sounds being covered; speech is clear and intelligible whenever it appears.

□ Play the downloaded file in full to confirm sound, picture, and timing; a successful preview does not replace export acceptance.

Review: Does the story still hold without the music? If not, check the narrative and picture first-don't just turn up the music.

Next chapter: Turn problems into an actionable revision list.

Further: turn the revision list into a film plan, see three scenes, twelve shots.

Advanced 6: Review and targeted fixes

Core question: What to fix first, which spot to change, and when can you stop redoing it?

Understand the principle first

Break "doesn't feel premium" into observable problems: you can't tell why the keychain matters, screen direction conflicts at a cut, a deformed finger hides the action, music covers the doorbell. Note the time and impact, then choose an operation. Fix delivery blockers such as unplayable files and missing assets first, then story confusion and key action errors, then continuity and sound, and only then local polish.

For each item, state the expected improvement and the smallest change: if you can trim, don't regenerate first; if a suitable candidate already exists, don't blindly reshoot. When the story arrangement is wrong, upscaling, frame interpolation, or switching models won't solve it for you. Set a revision scope and acceptable cost for this round, stop after comparing improvements, and don't treat endless retries as a quality guarantee.

Good vs. bad: written revision list

Turn aesthetic judgment into checkable actions.

Not actionable	Verifiable
"The whole film doesn't look good; start over."	"In S04 the hand deforms and covers the prop when it touches; try trimming the tail first, and redo only S04 if the action is missing."
"Switching to the HD model will definitely be better."	"The new candidate makes the action clear and keeps identity unchanged; after adopting it, update only the specified edit instance, then re-export and check."

TimeGleam hands-on practice and exercises

Download file / current timeline → issue log → candidate from original node or local redo → check the new canvas version

1. On the first full pass, watch without stopping and write down your story understanding and obvious issues; on the second pass, locate the time, shot, and actual asset. If there is no real file, pre-review with a written storyboard and do not claim you checked image quality.

2. Pick only the three items that most affect viewing and record "evidence-impact-smallest fix-recheck method". First try local edits on a copy; when generation is needed, check the target, references, and authorization; do not automatically redo successful upstream nodes.

3. Compare the original and revised versions. After adopting a candidate, check whether the specified edit instance has been replaced; keep unselected instances and old exports. Play the newly downloaded file and record unresolved items truthfully.

Acceptance checklist

□ Each item has a time or shot position, evidence, and observable improvement; do not just record the buttons you clicked.

□ Know what is done locally and what needs a model task; do not repeatedly generate beyond authorization.

□ The original issue improves without introducing new issues in identity, references, pacing, or sound; list any remaining issues as incomplete.

Full course completion: six-item self-assessment

Score STORY, shot intent, continuity, light and color, pacing and sound, and delivery 0/1/2 each: 0 means an obvious issue affects understanding; 1 means it basically holds up but you can point out how to fix it; 2 means it is clearly expressed and you can explain the tradeoffs. This is a learning self-assessment, not an automatic TimeGleam score, industry certification, or release threshold; unplayable files, misused assets, or authorization issues cannot be offset by the total score.

Work / version / date:
Goal for this round:
STORY__ Shot__ Continuity__ Light and color__ Pacing and sound__ Delivery__ (each 0-2)
Issue location (time / shot / asset):
Actual evidence and impact on viewing:
Smallest revision plan / what to keep:
Whether model costs are involved; authorized scope:
Recheck result after revision:
Unfinished items / entry point for next round:
Stop condition: the three items in this round are resolved and no new key issues were added; otherwise record the reason and do not retry indefinitely.

Review: Did this improvement come from the story, the shot, or the edit? Use a specific edit point or frame to explain; do not judge progress by how many points you spent.

For operation review, see Free canvas full manual; for multi-episode management, see Workflow guide; for Agent continuation, see MCP in practice.

Complete film plan: three scenes, twelve shots

What this section solves: This expands the earlier case into a complete short film. It is not another button to learn: it walks the overall plan, shot list, subject, first frames, video, sound, and edit onto the canvas in order, then exports a film you can watch from start to finish.

Fix the structure before touching the canvas

"The Last Lamp" runs 58 seconds across three scenes and twelve shots: 9 seconds on the street, 38 in the shop, 11 on the way home. The first scene establishes the character and his situation, the second carries finding, hesitating, and taking, and the third lands the emotion on the two engraved words "home". Give each shot one main action, 4-7 seconds each, and keep half a second of handle on both sides when editing so the image is not cut away the moment it settles.

Twelve-shot list, ready to paste into a shot draft

Shot	Set	Duration	Image and action	Sound
S01	Rainy street corner	5s	Wide shot. Wet asphalt reflects the light; Lin Mo walks in from the right under a black umbrella as the shop sign's warm glow bleeds into the rain.	Rain, distant traffic
S02	Shop door	4s	Medium tracking shot. He closes the umbrella, shakes off the water, and pushes the door; a bell chimes and the rain dulls behind the glass.	Door bell, rain turning muffled
S03	Inside the shop	6s	Wide lateral move. Shelves, pendant lights, and an old radio as he walks slowly in along the counter.	Low radio hum, footsteps
S04	Counter	5s	High angle, slow push. A worn brass keychain lies on the wooden counter; the engraved word "home" catches the light.	Rain, wall clock
S05	Reaction	4s	Close shot. He stops, fingers hovering above the keychain without reaching for it yet.	Breathing, fabric rustle
S06	Hand detail	4s	Macro. A fingertip touches the brass, the pad brushing the engraving; patina and scratches are clearly visible.	Light metallic ring

Shot	Set	Duration	Image and action	Sound
S07	Flashback	5s	Warm old kitchen; a hand hangs the same keychain on a wall hook, the image trembling slightly.	Wall clock, very low piano note
S08	Back in the shop	4s	Medium close shot. He comes back to himself, looks up at the shopkeeper, and confirms with his eyes.	Rain returning to the foregr ound
S09	Exchange	5s	Two-shot, medium. The shopkeeper nods; he sets coins on the counter and picks up the keychain.	Coins on wood
S10	Leaving	5s	Medium shot. He pushes out through the door, the bell rings again, and the shop light reflects in the glass.	Door bell, rain rushing back
S11	Alley mou th	5s	Wide shot. He stops, closes the keychain into his palm, and looks up at a lit window in the distance.	Rain, a distant television
S12	Final imag e	6s	Close shot, slow pull-back. The keychain in his palm, the word "home" clear, the image fading down.	Rain fading, one closing note

Do this: one route from a blank canvas to a finished film

New blank canvas → three sections → overall plan and shot list → subject and prop → each first frame → each video shot → ambience → timeline → export

1. **Create the document and sections:** Home canvas creation → New blank canvas, name it " The Last Lamp · Film plan"; add three sections, "Scene 1 Street", "Scene 2 Shop", "Scene 3 Homeward". Sections only affect layout and finding nodes; they do not trigger generation or change document ownership.

2. **Place the overall plan and shot list:** Write the overall plan in a text node and paste the twelve rows into a shot draft node. Confirm the current shot table before adding references; confirming a draft calls no model and spends no points, it only pins down the shots you are about to make.

3. **Build the subject and the prop:** Put a primary image of Lin Mo plus two supporting images in a subject pack, with notes for the dark grey worn jacket and the black umbrella; give the keychain its own material node shot close enough to read the engraving and patina. References must actually be attached to the node: naming an asset in the prompt alone does not send the image to the model.

4. **Make the three key inserts first:** Generate the first frames of S04, S06, and S12 first and check whether the engraving reads and the brass looks like an old object. If these three fail, the film is just "someone picked something up"; once they pass, fill in the other nine and avoid reworking a whole batch.

5. **Only animate after the first frame passes:** Each video node references that shot's first frame and its storyboard prompt, 4-7 seconds per shot. Two candidates per shot is enough: judge action and camera movement first, image quality second. Keep one camera intention per shot; do not ask for a push, a move, and a follow in the same shot.

6. **Adopt explicitly instead of rebuilding in new nodes:** Preview the candidate and click Use to write it into the node; keep one formal result per shot and leave the rest in history. To redo, change only that shot's prompt or first frame and regenerate rather than duplicating nodes and rerunning the film.

7. **Build sound as its own layer:** Treat ambience, spot effects, and the flashback score separately: rain as a bed, the bell, coins, and metallic taps aligned to their picture hits, and the flashback piano kept very low as texture. The no-dialogue cut can skip voice entirely; adding it tends to steal the scene.

8. **Assemble by shot number:** Drag S01-S12 onto the timeline in order, rough-cut for rhythm first, then handle transitions. Use a short dissolve into the S07 flashback and hard cuts elsewhere; cut in the middle of an action, never on the stiff pose just after it ends.

9. **Export and watch it all the way through:** Export 1080p and play the downloaded file from the start: check duration, the engraving, how the bell meets the rain, and the light change after S10 leaves the shop. A playable preview is not a successful export; only a downloaded file that plays counts as done.

What each canvas detail is for

The twelve shots across three scenes, and where they most often break

Where	What to do	Why it matters	How to check
Document and sections	Three sections for three scenes, nodes named by shot number	You will not hunt across the canvas later	Searching "S07" jumps straight to it
Subject pack	Primary image, two supporting images, wardrobe notes	Faces and clothing drift most across scenes	Pull one frame per scene side by side; face and jacket match

Where	What to do	Why it matters	How to check
Reference chips	First frame and keychain images actually attached to nodes	A name in the prompt sends no image	The node's More settings shows the reference thumbnails
Models and specs	16:9, 4-7 seconds per shot, 1080p export	Specs drive both the quote and the film's aspect	The quote page matches what you actually set
Shot draft	Twelve rows, one shot and one action per row	Let the model do one thing at a time	One main action per row
Candidates and adoption	Two candidates per shot, one adopted explicitly	Generating a candidate does not adopt it	The node carries the current-use marker
Audio nodes	Ambience, spot effects, and flashback score on separate layers	Disconnected sound breaks the spell fastest	Listen once with the picture off
Timeline	Shot order, mostly hard cuts	Order carries the story	It plays start to finish without a break
History and exports	Keep the old export, add a new one	You can fall back to the previous version	History holds two or more exports

One prompt for the whole plan

"The Last Lamp", a complete short film, about 58 seconds, 16:9, 1080p export, realistic cinematic look.
Three scenes, twelve shots:
S01 Rainy street corner 5s Wide shot, Lin Mo walks in from the right under a black umbrella, wet asphalt reflecting light, the shop sign's warm glow bleeding into the rain.
S02 Shop door 4s Medium tracking shot, closes the umbrella, shakes off water, pushes the door, a bell chimes.
S03 Inside the shop 6s Wide lateral move, shelves, pendant lights, an old radio, he walks in along the counter.
S04 Counter 5s High angle slow push, a worn brass keychain on the wooden counter, engraved with the word "home".
S05 Reaction 4s Close shot, he stops, fingers hovering above the keychain.
S06 Hand detail 4s Macro, a fingertip brushes the engraving, patina and scratches clearly visible.
S07 Flashback 5s Warm old kitchen, a hand hangs the same keychain on a wall hook.

S08 Back in the shop 4s Medium close shot, he comes back to himself and looks up at the shopkeeper.
S09 Exchange 5s Two-shot medium, the shopkeeper nods, he sets coins down and picks up the keychain.
S10 Leaving 5s Medium shot, he pushes out through the door, the bell rings again.
S11 Alley mouth 5s Wide shot, he closes the keychain into his palm and looks up at a lit window in the distance.
S12 Final image 6s Close shot slow pull-back, the keychain in his palm with the word "home", the image fading down.
Shared constraints: the same character (Lin Mo, 25, worn dark grey jacket), the same old grocery store, the same worn brass keychain engraved with the word "home"; warm shop light against a deep blue rainy night; no dialogue, no subtitles; natural ambience as the base, with a very low piano note possible in the flashback.
Before generating, give me the shot list and the first-frame plan for every shot, and wait for me to confirm scope and budget; produce each first frame first and only animate after I confirm; show candidate comparisons before adopting anything and never overwrite existing results.

Estimating budget and time

Planning, twelve first frames, twelve video clips, and optional audio are priced separately; local editing and export add no model fees, and actual points follow the quote page. Generate the plan without confirming first, read the two separate lines for settled planning and this hold, then decide between a three-shot test and a full run. Reworking a first frame costs far less than reworking video, so spend the extra time on S04, S06, and S12.

Acceptance checklist

□ The three scenes read apart at a glance: cold blue rain, warm yellow interior, cold blue on the way home with a hint of window light.

□ No action repeats: the fingertip reaches the keychain in S06, and only then does S07 enter the flashback.

□ The engraved word "home" is clearly readable in S04, S06, and S12, not smeared.

□ With the picture off, the bell, coins, and rain still connect.

□ The downloaded file plays from the start and lands within two seconds of the shot list.

Review: If you could keep only three shots, which three? Do those three make it clear why he takes it?

Go deeper: Use the review sheet to locate problems, for pricing and independent editing details see One-click film handbook, to hand it to an Agent, see MCP in practice.

06 Desktop app, MCP, and automation
plug a local Agent into the same creative space

TimeGleam Desktop: put an Agent workbench in the same creative space

The web studio owns the canvas, timeline and finished film; the desktop app is a second workbench installed on your machine that plugs a local Agent, one that can read and write files and run commands, into the same TimeGleam space. It is not a phone-sized web page; it is the execution end that can touch local files and still call remote tools. Learn the split first: the web does picture and edit, the desktop does orchestration, batch work and delivery.

Download the installer: TimeGleam Desktop for Windows (Windows 10/11 x64). The page lists a fixed download URL and the SHA-256 of the current installer so you can check it after downloading. The installer is not code-signed, so the first run shows Windows SmartScreen's "Windows protected your PC"; choose More info, then Run anyway. That is normal for an unsigned build rather than a damaged file. If the hash does not match, do not run it.

What the desktop app is, and how it differs from the web

One account, one space, two entrances

Capability	Web studio	TimeGleam desktop app
Canvas, timeline, export	Complete	Opens the same set of pages in a browser
Local file read and write, batch renaming, scripts	Limited by the browser	Native access inside folders you authorise
Calling remote tools (MCP)	Not directly	Built-in connection, available after sign-in

Capability	Web studio	TimeGleam desktop app
Keeps running without the browser	Needs the page open	Keeps working in the background after you close the window
When to use it	Look at frames, change an edit, check quotes and order	Organise assets, batch-submit, write scripts, check jobs

There are not two copies of the data: the desktop edits the same projects, episodes, canvases and films in the same space. Adoptions and timeline changes made on the desktop appear on the web after a refresh, and documents saved on the web are what the desktop reads at the latest revision.

First launch: sign in, pick a space, prove the connection

1. **Install and launch for the first time:** Get the Windows installer from the download page and double-click it; the first launch stops at the welcome page and asks for an access key. The installer is not code-signed, so the first run needs More info, then Run anyway in the SmartScreen prompt. To upgrade on the same machine, run the new installer over the old one: sign-in and settings are kept. Different product lines install separately and never overwrite each other.

2. Sign in through the browser Click Sign in on the welcome page and the browser opens the TimeGleam account page; once you confirm, the desktop app finishes signing in, with no key to copy by hand. You can also paste an existing tmg_ key.

3. **Confirm the workspace:** After signing in, check that the current space is the one you want to create in. If the account has no space, the welcome page will not let you through; create one on the web first, then come back.

4. **Run one minimal check:** Have the Agent list the projects once. Real project names and IDs mean the connection works. Do not start by asking it to produce a finished film.

> **Key lifetime and revocation:** Desktop keys are valid for one year and renew automatically while you use the app, up to two years; after two years, confirm once more in the browser. Changing machines, sharing a computer or suspecting a leak all call for revoking the key in Account Security on the web under connected tools, then signing in again. Never paste a key into a script, a prompt or a chat log.

Two things the desktop app really has

The first is **remote tool connection**: after sign-in the desktop connects to TimeGleam's remote tool service, so the Agent can list projects, read episodes, edit canvases, check jobs, get quotes and export. These map one to one onto the web buttons; there is no desktop-only hidden generation endpoint. The second is **local execution**: the Agent can read and write files, run scripts and organise assets inside folders you authorise, which the web cannot do.

The desktop also ships two local image tools: one lists the image models available to the current space, the other renders with a model and saves the result as a local attachment. They bill to the space the same way web rendering does, with quotas computed on the server; the desktop does not set its own prices. Scripts, review and organising are free to try; the moment you render images, generate video or make a film, the quote-then-confirm rule still applies.

Where the desktop app fits in the process

The recommended order: fix the story and specs on the web → organise assets and batch-submit on the desktop → review candidates, adopt and cut on the web → organise delivery files and receipts on the desktop. Judging picture quality needs a large view and real playback, which the web does best; repetitive work (renaming, aligning shot numbers, collecting receipts, checking statuses in bulk) is fastest on the desktop.

The desktop can also make the film outright: the Agent queries the project, reads the plan, gets a quote, confirms, waits for the job and fetches the download links. That suits people who already know the process and just want it run. The first time through, still do it on the web so you know what each step should show, then hand it over.

Acceptance checklist

□ You can name the account and space you are signed in to, and know the desktop and the web see the same content.

□ The Agent returns real project IDs rather than plausible-looking ones.

□ You know where to revoke a key and that changing machines means signing in again.

□ You know the desktop offers no way around the quote.

Review: If you could keep only one entrance, which would you use to judge and which to execute? Being able to say why means the split is clear.

Next: use Don't Knock drill to walk the whole path from story to finished film, or MCP in practice for tool-level detail.

MCP hands-on: continue in the same work

Remote MCP lets AI tools use saved works, assets and tasks in TimeGleam; the web app and agents can continue the same project. First finish one shot in the eight lessons, then hand repeatable, local, checkable steps to external tools.

> External tools are not a remote mirror of every web feature. Remote MCP now covers canvas final cuts, single-shot continuation and basic editing; it still does not expose the professional image workbench, frame-by-frame shot breakdown, professional audio processing, or director scenes; do not invent corresponding tools or have agents bypass the web app's permanent cleanup confirmation. The REPLACE IDs, versions, models, durations and parameters in the examples below must be replaced with real responses and cannot be treated directly as business facts.

All examples teach structure and do not run automatically. Practice in a test work before writing; generation and recovery may consume credits, so obtain authorization first. You can download offline Teaching JSON examples and All tool schemas in this version; the PDF also contains these two JSON files. PDF readers that support attachments can extract them; if they do not, use the web download. For online operations, follow the MCP inputSchema returned by the service.

Connect an AI tool

1. In TimeGleam, go to "Account security → Connect an AI tool", choose a client, and copy the service URL.

2. In the client, add remote MCP and use OAuth to register automatically; sign in on the TimeGleam page that opens and confirm the workspace and authorization purpose.

3. Return to the client and first read identity_get and creation_guide. Read-only members can only view; generation requires current edit permission, model authorization and cost confirmation.

4. You can revoke anytime in TimeGleam "Connected tools". Switching workspaces on the web does not change existing authorization; to switch the authorized workspace, reconnect.

No Timegleam-specific program to install and no key to create or save by hand. Client plans, organisation policies and entries are set by their vendors; Timegleam cannot remove those limits. Real-client compatibility for this version still needs acceptance; protocol tests are not counted as measured support.

MCP: First understand IDs, revisions and receipts

The most confusing IDs

Field	Actual meaning and source
production_id	Work ID, read from real responses such as production_list.
project_id	The current document's scope.documentProjectId; a workflow is a formal episode, and a free canvas is an independent document, not a work name or displayed episode number.
canvas_id	Section within the document, defaults to main; not the free document's project_id.
shot_id / node_id	shot uid and the real node ID, read from the saved node or the shot's imageNode/videoNode; never build an ID from "S04".
revision / production_revision	Document revision and shared work revision; read before each edit, and check both.
request_id / submission_id	A stable ID for the user's intent in this attempt, 8-64 alphanumeric characters, underscores, or hyphens; if the network drops the response, keep the original ID and original body.
impact_token	The 64-character credential returned by the corresponding preview; confirm it together with the original selection, body, and version, and do not fabricate it.

Save in the web app first, then read with episode_get. Regular canvas commands, final adoption, candidate adoption, edit changes, media sync, and export follow their respective preview flows. A preview not writing does not mean a read-only scope is always callable; service permission checks still apply. Create, upload, and finished-film writes use request receipts; after confirmation, query the original film_id and task. Do not treat every endpoint as the same kind of idempotent retry.

> First read identity_get, production_list, and creation_guide, then call tools using the inputSchema returned by the service.

MCP returns structured ok and data; failures include error.code, message, status, and isError. Read the actual receipt; do not treat a chat description as saved.

Walkthrough 1: Start in the web app, edit with MCP, continue in the web app

1. Open "The Last Lamp" in the web app and save it. Use production_list, episode_list, and canvas_list to locate the actual document; use episode_get to verify scope, nodes, and both revisions. If there is no work, use creation_draft_create to create a draft, a first episode, and a free canvas without generating.

2. Prepare the JSON below, first replacing project_id and the current revision; request_id represents this intent, so choose an unused number for the first execution. Call canvas_command_preview and inspect created_nodes, impact, and impact_token.

3. Copy the original JSON, add only the real impact_token and call canvas_command_apply. Do not change the command, the IDs or the revision. Reload the page and you will see two ordinary nodes, "Rainy-night idea" and "S04".

4. Prefer media_import to receive host attachments, or have a client with file permissions upload via media_upload_prepare; if unsupported, browser_open returns the project upload page. When done, use asset_list to read back the stable asset ID.

```
{
  "project_id": "REPLACE-DOCUMENT-ID",
  "request_id": "lamp-create-nodes-001",
  "revision": 1, "production_revision": 1,
  "command": {"operation": "create", "nodes": [
    {"key": "idea", "kind": "text", "label": "Rainy-night idea",
     "text": "Lin Mo finds a brass keychain engraved with 回家 in a rainy-night grocery store."},
    {"key": "s04", "kind": "video", "label": "S04",
     "prompt": "Lin Mo's gaze drops to the counter as the camera pushes in slowly."}
  ]}
}
```

```
Call the canvas_command_preview tool; for parameters, see canvas-create.json
# Review the preview manually, add impact_token to the original file, and save it as
canvas-apply.json
Call the canvas_command_apply tool; for parameters, see canvas-apply.json
Call the episode_get tool; for parameters, see document.json
```

The ten partial commands each submit only their corresponding fields.

operation	The content required by command (you may also provide canvas_id)
create	nodes: up to 20; temporary keys are only for this batch; read subsequent IDs from the preview's created_nodes.
update	node_id, patch; changes only regular objects; the parameters field is replaced in full.
copy / delete	node_ids; copy/remove regular objects; the original media is not deleted.
connect / disconnect	edges (source, target, role) / edge_ids; preserve explicit order and do not duplicate edges with the same source and target.
group / ungroup	node_ids and name/group_id; do not modify managed production objects.
rename	name; only renames a free document, not a work or a formal episode.
references	node_id and references; each item is only one of asset_id or node_id, with role set to reference/first/last.

You can create six types: text, image, video, audio, storyboard, and reference; for a subject, provide the name, description, primary image, up to 11 supporting images, and voice through the reference's subject field, and do not fabricate a formal identity. Create a reference for existing files; do not write an assetId directly to a generation node to pass it off as a successful result. Regular commands do not support whole-document overwrite; formal, locked, and managed objects continue to use dedicated entry points.

canvas_create uses a formal episode ID to create a free document; source_id can copy a free document from the same episode; canvas_recycle soft-recycles against the latest dual revisions, returns not_found on repeats, and does not restore or revive. subject_list only reads existing asset packages and does not bind automatically. For a cross-client 409, first read the current state, compare, then preview the new intent; do not just replace the version number and force an overwrite.

Hands-on 2: Check inputs and costs, and edit only the target shot

1. First use project_context_get to verify the project, canvas, and production ID, then use film_get to read the target film. model_list and film_options_get read the currently allowed models and specifications.

2. If a complete plan already exists, save it with film_create, reuse stable asset IDs, and do not call a planning model separately. For generated images, use project_images_get to view thumbnails by node; media_delivery can deliver limited previews and direct downloads, but the first frame cannot replace a full review.

3. Before modifying an existing film, call film_pause and wait for in-flight tasks to finish. Use film_shot_preview / film_shot_apply to verify shot_key and the impact; keep old results and do not redo other shots.

4. After updating the affected steps, verify pricing with film_quote. Show the model, shots, duration, specifications, credits, and limits in quote_summary, obtain the user's explicit consent, then use film_confirm to submit delegated, the original quote fingerprint, and the quota. The web confirmation_path remains a compatible entry point; returning only pending confirmation does not mean the task has started.

5. After confirmation, use film_get to query the original production and background task. If disconnected, continue querying the original film_id; do not create another production.

```
{
  "project_id": "REPLACE-DOCUMENT-ID", "request_id": "lamp-reference-001",
  "revision": 2, "production_revision": 1,
  "command": {"operation":"references", "node_id":"REPLACE-VIDEO-NODE-ID",
    "references":[{"asset_id":"REPLACE-IMAGE-ASSET-ID","role":"first"}]}
}
```

```
{
  "project_id":"REPLACE-DOCUMENT-ID", "film_id":"REPLACE-FILM-ID",
  "request_id":"lamp-quote-001", "revision":3
}
```

Cost confirmation is valid only for the current quote and version; re-verify after the quote, model specifications, or content change. Explicit consent in the session allows an authenticated client to submit on the user's behalf, and web confirmation is still available; do not treat connection authorization as unlimited spending authorization.

Regular canvas references can still be previewed and saved; when generation is needed, go through the film flow with quote confirmation, or have the user act on the web.

Continue the same film across clients

Use work_search / work_entries_get to locate this workspace's entry, and film_list / film_get to read the film_id and shot bindings. If a plan or storyboard already exists, use film_create and do not repeat model planning. film_quote saves the quote; after the user explicitly consents in the session, film_confirm starts production with delegated, the original quote fingerprint, and the quota; web sign-in confirmation remains available.

For a single shot, first film_pause and wait for in-flight work to finish; use film_shot_preview / film_shot_apply to verify shot_key, the finished film and document versions; after saving, update the affected steps and re-confirm pricing. Old assets and other shots are retained. film_editor_preview / film_editor_apply only change the specified finished film's basic edit; film_export_preview / film_export_submit export again; use film_get to continue checking, and receipt_get for the original request. In the web page, save first, then reload; after an old preview returns 409, compare and preview again. The following timeline_* examples are still historical project timelines and do not apply to film_id.

Prefer stable asset_id for assets. When the host provides attachments, use media_import; clients with local file permissions use media_upload_prepare to upload raw bytes, with a link valid for five minutes. If neither is available, only open the current project's upload page, and the Agent reads back the result afterward. For delivery use media_delivery, read the actual file specs and download directly; 480p can run through production and export, and do not upscale 720p to pass it off as 480p. session_capabilities can only report capabilities actually declared by the host; do not assume Codex has ChatGPT attachments or a player.

Hands-on 3: Assemble an edit, download, and rework a section

First use timeline_get to read the currently saved timeline. The example below adds only one video track and one clip, keeping old tracks; the duration must be within the actual source video range. Call timeline_preview to verify, then timeline_apply with the original request_id and command plus the real impact_token.

```
{
  "project_id":"REPLACE-DOCUMENT-ID",
  "request_id":"lamp-timeline-001",
  "command":{
    "newTracks":[{"key":"lamp-main","name":"Video","type":"video"}],
    "additions":[{"trackId":"lamp-main","type":"video",
      "name":"S04","start":0,"duration":4,
      "assetId":"REPLACE-VIDEO-ASSET-ID"}]
```

```
  }
}
```

changes adjusts start, duration, sourceIn, playbackRate (0.25-4), volume (0-2), text, name, asset, track or remove by the real elementId; splitAt uses absolute time and is submitted alone, without other changes mixed in. newTracks allows at most 8; changes and additions allow at most 20 items each. trackChanges controls mute/hide. Basic editing is not the full professional web editor; do not include effect fields outside the schema.

1. After timeline_apply succeeds, verify the clips, audio/video and save status in the web page. export_preview reads the currently saved snapshot and checks duration, dimensions and missing assets.

2. export_submit carries the original request_id and real impact_token; record the persistent job_id, and only after the original task succeeds take the official asset_id from the result. It is local composition and does not call paid models.

3. browser_open returns the entry points for viewing and downloading on the web page.

4. To redo shot 3, modify only the target with shot_update or a normal node update, then generate and adopt only that node. timeline_media_scan finds old instances; timeline_media_preview selects elementId and nodeId; after confirming, timeline_media_apply. Preserve timing and compatible manual effects, and do not update unselected instances.

```
{
  "project_id":"REPLACE-DOCUMENT-ID",
  "request_id":"lamp-sync03-001",
  "selections":[{"elementId":"REPLACE-MCPP-ID",
                 "nodeId":"REPLACE-VIDEO-NODE-ID"}]
}
```

Acceptance: The web app and MCP read the same timeline and change only the target segment; the new export ID coexists with the old export. Delivery is complete only when the file downloads successfully and the video plays. Repeating the same export request and snapshot returns the original job; do not create another export just because the network is waiting.

Walkthrough 4: Adopt free content into the official workflow

1. workflow_setup_get reads the official episode setup items, and workflow_enable confirms the name and style against the latest dual revision. Enabling does not generate content or automatically turn free content into official content.

2. Call adoption_targets_get from the original free document, select only valid targets that belong to the episode, and do not guess IDs across episodes or workspaces.

3. adoption_preview takes the real source_node_id and kind: visual can create a new item or a new version on card_id; shot can create a new item or target shot_id; storyboard uses row_ids to select draft rows; script explicitly sets append or replace.

4. Review the frozen source, target version, and materials, add the real impact_token to the original selection, then call adoption_apply. If the source/target revision or material bytes change, redo the preview.

5. If there is no response, query adoption_receipt_get with the original free document project_id and request_id. Replaying with the original ID and original body does not adopt twice; the source and old version are kept, and no clip is automatically generated or replaced.

```
{
  "project_id":"REPLACE-FREE-DOCUMENT-ID",
  "request_id":"lamp-script-adopt-001",
  "source_node_id":"REPLACE-TEXT-NODE-ID",
  "kind":"script", "script_mode":"append"
}
```

Having an image does not mean you can lock an official identity: visual lock_reference must be selected explicitly, and status versions also need a valid parent_version_id; do not set all subject reference images and audio as official primary references. storyboard supports at most 20 rows at a time; keep the original row IDs and do not treat display numbers as row_ids. Verify official shots and script reads with production_get, script_list, episode_get, and similar.

Acceptance: The official script in the web app shows the added paragraph or explicitly replaced content, and the free text remains. Undoing the chat does not undo a successful adoption; for later changes, start a new real intent instead of editing the old receipt.

Walkthrough 5: Lost responses, 409, and interrupted jobs

Case	Safe continuation
Lost response on create/upload/ordinary command	First, use receipt_get to check the original request_id. A 404 may mean the request is still running, was not submitted, or was deleted, so do not blindly change the ID; if needed, replay with the original ID and original body. A deleted object will not come back.
Lost response on generation submission	Check job_list against the original record; replaying with the same submission_id and the exact same body returns the original job, while changing the body returns 409.
Lost response on official adoption	Call adoption_receipt_get in the original source document; do not switch to the target episode ID, and do not create another adoption.
409 on update	Stop the old write, reread with episode_get/timeline_get, compare later edits in the web app, and preview the new intent again. Do not just change the version number to overwrite someone else.
watch timeout	Stop only the local wait; the server-side job continues. On a different client, query with the same workspace identity and the original job_id.
Job interrupted/failed	First verify with job_get and film_get; for a finished film, resume only with authorization under a valid quote; handle regular planning/generation tasks in the TimeGleam task center. job_cancel does not guarantee the provider won't charge.
401 / 403 / 404 / 422	Verify credential validity, scope and member permissions, object visibility, and input schema separately; do not output tokens or switch to higher-privilege credentials to bypass restrictions.

After a network interruption, first check the original task and receipt, then continue with the original ID and original body. On version conflicts, reread and preview again; do not automatically resend write requests and overwrite web edits.

```
creation_guide (workflow=resume)
Call tool receipt_get; parameters in receipt.json
# receipt.json: {"request_id":"lamp-create-nodes-001"}
Call tool job_get; parameters in job.json
# job.json: {"job_id":"REPLACE-JOB-ID"}
Use job_get to query the original job
```

Offline drill: save the original request, ID, response, and task ID (excluding credentials), and walk through the three lost-response scenarios verbally. You do not need to actually interrupt a paid task to practice; validation in this tutorial uses simulated transport only and does not claim that real provider recovery was performed.

MCP integration, three guides, and Agent prompts

First complete browser authorization from "Connect an AI tool". creation_guide provides new-video (create from scratch), retake (redo a single shot only), and resume (continue after disconnection). Supported clients can also read the same-named Prompts and timegleam://guides/ resources; when these entry points are not shown, the basic tool flow is still available.

For local assets, browser_open returns the current project web page; after the user uploads them, read back stable asset IDs with asset_list. Preview and download finished films through the TimeGleam web page; do not pass server file paths, scrape arbitrary URLs, or use permanent public links to bypass sign-in. Reading guides does not generate or consume credits; generation requires price verification and confirmation.

Use the connected TimeGleam MCP. First read identity_get, model_list and creation_guide(retake), and locate the real space, production, episode and free document for The Last Lamp. Read the saved version and tell me the actual node ID of S04, the current take and the timeline reference. Offer only the local change and a cost range, then wait for my authorisation before previewing and applying; keep the other shots and the old media. Check the real inputs before generating and record the original submission_id and job_id; after a lost response check the original job instead of paying twice. After adopting a new candidate, preview only the timeline instance to be replaced, then export to a new file. Do not call retake, director or permanent-cleanup interfaces that are not open yet.

Completion criteria: the client can list tools and the three guides, identity binding is correct, and read-only requests succeed; state separately whether writes and generation are authorized. Authorization credentials must not enter shared tutorials, project packages, or Git.

07 Freeform Canvas Manual
The complete manual, organised by feature

Complete feature manual for the free canvas

The eight lessons teach one shortest creation path; this manual is for looking things up when a concrete task appears. The following follows version b3a7988c (2026-10-19), covers currently reachable entries and does not mean the running service has been upgraded. All exercises reuse The Last Lamp with copies and your own footage; generation is billed only when you explicitly submit a model job.

Four kinds of "Adopt" cannot be used interchangeably: a text candidate replaces/appends to the body; a media candidate becomes the node's current result; free content is explicitly adopted into the formal workflow; the timeline is checked separately and the specified clip updated. A successful step never automatically completes the other three.

Canvas manual: documents and nine object types

Top: Workspace / Work / Episode → Mode and Canvas → Manage canvases; bottom: Add or double-click empty space

Each free canvas has independent content, saved revisions, and tasks, and belongs to one episode of a work. Free canvases in the same episode share the work's creative bible and visual identity, but a regular save does not overwrite the shared settings. The workflow canvas is another view of the formal episode, not a separate free document that can be deleted. Manage canvases lets you create, rename, duplicate, and soft-delete; duplication keeps media but does not reuse the old task execution identity.

Choose the right object instead of cramming everything into an image node

Object	Best for
Film	Holds one-click creation records, film previews, standalone edits, and export history; it is not a regular video generation node.
Text	Briefs, concepts, body text, and text candidates; can feed into shot drafts.
Image / Video / Audio	Configure the corresponding model and inputs, then submit generation manually; standalone audio is mainly text voiceover.

Object	Best for
Shot draft	Edit per-shot scenes, actions, references, and prompts, then choose shot generation or formal adoption.
Reference materials	Holds existing files and their sources; uploading files does not require configuring a model.
Subject / Visual asset	Organizes character and other asset packs and reads formal identity versions; a subject does not automatically create a formal identity.
Frame-by-frame breakdown	Analyzes the source video's shots, motion, and music, and outputs grouped references; it does not generate new video.

Exercises and acceptance

Create a new "Lighting Test" Free canvas, add text and a reference image, save, switch back to the original canvas, then return. The two documents' contents are independent; the work and episode at the top remain unchanged. Resuming recent work should restore the original document; do not repeatedly create new works to find nodes.

Finding works: My works and Works in this workspace

Home top left → My works; bottom of home → Works in this workspace; work cover → select a specific creation entry

First answer "which workspace, which work, which episode, which document," then continue creating. Both lists show only non-deleted works you currently have access to; they are not asset directories or lists of exported videos.

Three entry points, three search scopes

Entry point	When to use
Works in this workspace	Searches only the current creation workspace. Search by work name and paginate; most recently saved comes first by default. Browsing or task heartbeats do not count as content changes.
My works	Find works across all accessible workspaces, with workspace filters, name search, pagination, and grid/list toggle. Browsing and filtering do not switch the current creation workspace; only actually opening a target enters its workspace.
My assets	Find images, videos, audio, subtitles, and character setup; use it to reuse individual assets, not as a substitute for work and episode navigation.

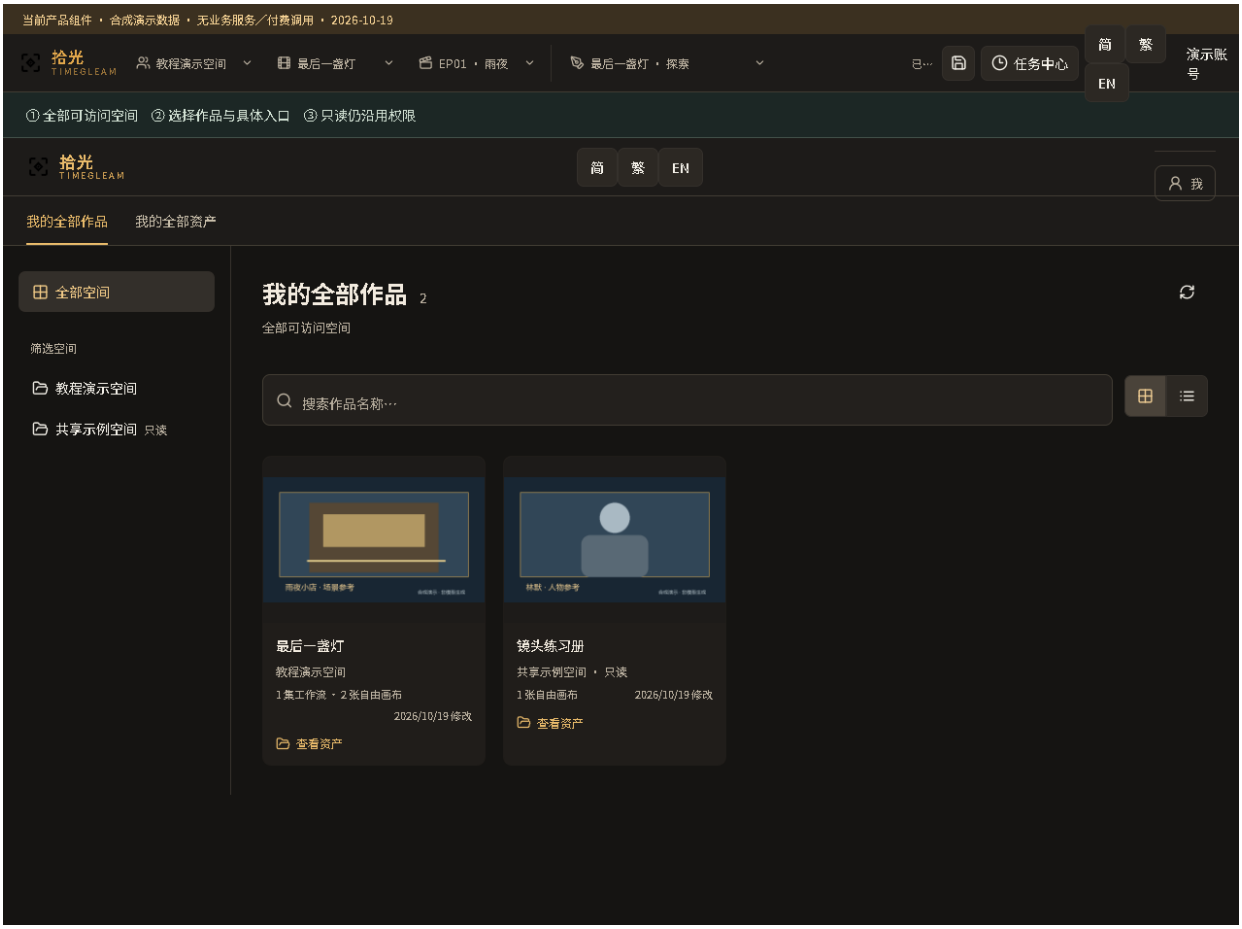
Enter the correct document from the cover

1. In "My works", first check the card's workspace, name, cover image, and last modified time; read-only workspaces are labeled read-only, and the aggregate view does not reassign team works as personal works.

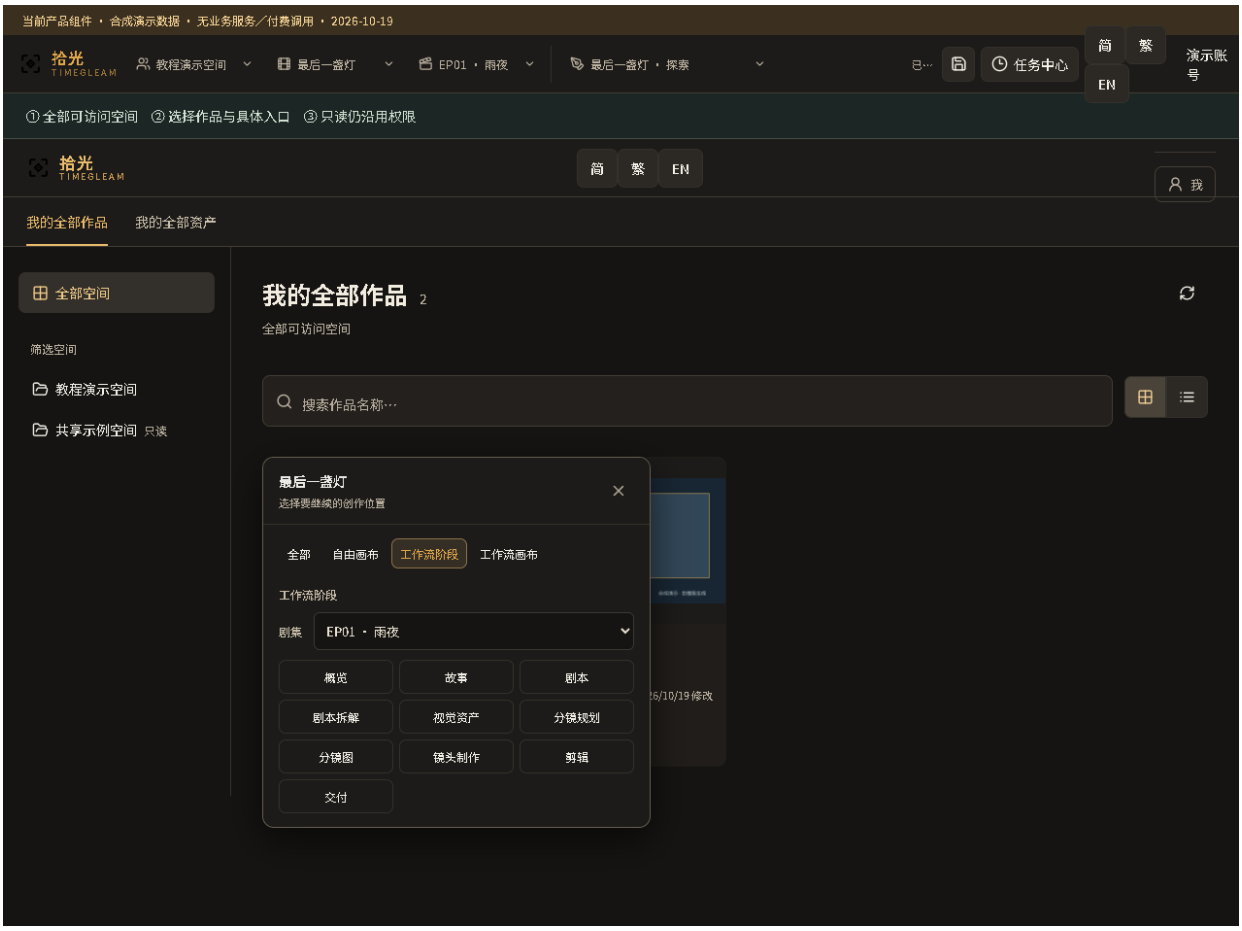
2. On desktop, hover over or focus the cover to view entry points; clicking the cover also expands them. On mobile, tap to show the entry panel. Covers can use an image or a cached video first frame; when preview is unavailable, a placeholder appears-this does not mean the work is lost.

3. Under "Free canvas", choose a specific document; under "Workflow stage", choose an episode first, then a stage; "Workflow canvas" opens that episode's main canvas. Do not assume clicking a card always enters the first episode. Draft workflows that are not yet enabled do not appear as entry points.

4. Read-only members can search and open, but subsequent actions are still limited by the original page permissions. If you cannot find something, first clear the search and switch workspaces, then check whether it was deleted or permissions changed; if loading fails, use " Reload" and do not create a work with the same name to recover old content.



The cross-space summary keeps the original space ownership; the example includes a read-only shared space. 2026-10-19, current components with composited data, not a user's real works.



Choose a document or episode stage before opening; expanding the panel does not generate or create content. 2026-10-19, composited work entries; demo callbacks do not perform real cross-space navigation.

Exercise: without creating a new work, locate "The Last Lamp" from "My works" and identify the two Free canvases "Exploration" and "Lighting Test"; then confirm the Workflow stage and main canvas entry points for EP01. Read the permission prompts; do not generate or delete anything.

Asset reuse: My assets and cross-workspace management

Home top left → My assets; asset panel in a work → My assets; the original work's asset panel remains available

"My assets" aggregates work assets from all workspaces accessible to the current account; browsing does not switch the workspace you are creating in, and files still belong to their original workspace and work. Read-only workspaces can be browsed and downloaded from; uploading, modifying, trashing, and copy destinations require edit permission.

Filter the scope first, then confirm where it will be used

1. Filter by workspace, work, image/video/audio/subtitle, and creative use; search materials, works, or descriptions; switch between grid/list as needed, paginate, and multi-select across pages. View characters and settings separately; they cannot be treated as ordinary image categories.

2. Open asset details to view preview, name, use, description, specifications, and "Where used"; "Go to source" returns to the original content. "Original file" downloads a single file. Only save name, use, and description when you have edit permission; when a concurrent edit notice appears, reread and verify, and do not overwrite someone else's changes.

3. "Upload materials": first select the target workspace, work, and the episode/canvas they belong to and use, then select files. Do not click model generation just to upload for archiving.

4. After selecting, choose a reuse method based on the differences below, and verify the target and result receipt; if a batch partially fails, review the reason for each item, handle only the failed items, and do not treat a success notice as full completion.

Copying is not the same as referencing, nor the same as generating.

Action	Results and boundaries
Copy to work	Create independent material and file in the target, with same names automatically distinguished; the original material is not moved, and later changes to the source file are not automatically synced to the copy.
Use for creation	First copy to the target work, then add a reference node on the selected free canvas; this does not submit a generation or automatically apply to the formal workflow.

Action	Results and boundaries
Copy character and settings	Copy the selected version, necessary base settings, and reference materials to create an independent draft version in the target; it does not include original shot bindings, historical candidates, or vendor identity registration. To edit the full settings, still return to the source work.



Check the space, work and material use first, then copy or use it for creation. 2026-10-19, current asset-page components with composited footage; no business data written and no model called.

Download and recycle boundaries

Media materials can be "packaged for download": up to 200 items and 2 GiB each time; " Download tasks" shows progress and errors; packaging continues after you close the page, completed packages are kept for 24 hours, and up to 10 tasks are retained. If partially completed, check the failure list first; if expired or interrupted by a service restart, package again. The package includes a directory and material list but is not a complete work backup; character settings themselves do not currently support packaging, and original files can be downloaded individually.

"Material Recycle Bin" provides soft-delete and restore, not permanent deletion. Before recycling, check references in canvases, settings, edits, and in-progress tasks; if there are

references, handle them at the source first, and do not automatically unbind or force-delete.Recycled work or episode materials do not appear in normal directories; permanent cleanup isperformed only according to Cleanup guide workspace permissions and web impactconfirmation.

Exercise: Filter Lin Mo reference images from "The Last Lamp" and view where they are used; in your own test work, use "Use for creation" and verify the copy's ownership and canvas reference node. Only inspect the result and do not click Generate; do not use actual recycling or deletion to verify the tutorial.

Canvas guide: navigation, selection, and shortcuts

Bottom canvas tools/canvas navigation; a selection action bar appears after multi-selecting nodes

The selection tool selects nodes and connections, and the hand tool pans; hold Space to pan temporarily. The tool group can be dragged and collapsed; collapsing does not remove functionality. The minimap locates positions on a large canvas; hiding connections only changes the display, and snapping only changes the layout, without canceling real references. Search for a node name and click a result to locate it; "Canvas content" can also search by body text and prompts.

Click the canvas before using shortcuts; while you are typing text, editing shortcuts are not captured; on Mac, use □ instead of Ctrl.

Key	Action and boundaries
V/H/Space	Select / Hand / Temporary pan.
Tab; double-click empty space	Open the add menu when the canvas is focused.
Ctrl+F / Ctrl+A	Search nodes / Select all nodes.
Ctrl+D	Duplicate selected regular nodes with their input references; the action bar also offers a copy without inputs.

Key	Action and boundaries
Delete / Backspace	Remove selected regular nodes or connections while keeping the media; this is undoable. Production objects use their dedicated management entry.
Ctrl+Z; Ctrl+Shift+Z / Ctrl+Y	Undo / redo canvas actions; does not reverse submitted model charges.
Shift+1 / Shift+2	Fit all / selected nodes.
Ctrl+0 / Ctrl+1; Ctrl+plus / minus	Fit all / 100%; zoom in / zoom out.
Alt+Shift+F	Auto-arrange; preserves content and generation relationships.
Ctrl+/; Esc	Expand / collapse assistant; deselect or close the current overlay.

Box-select multiple nodes to locate, group, duplicate, rearrange, or clear the selection. You can drag from node title bar text and blank areas; buttons still perform their own actions. Color markers are for classification and locating only and do not trigger generation. Read-only users can search by color but cannot change markers. Panel dragging and restoring the default layout are saved per user in this browser and do not move nodes. Practice: duplicate a set of references and video drafts, try one copy with inputs and one without, and check the connection differences; do not submit a generation just to test shortcuts.

Canvas guide: sections, groups, content, and templates

Top More tools → Sections and groups / Canvas content / Workflow templates / Canvas settings / Work creation settings

Sections within a document are not separate free canvases

In "Sections and groups," create or rename sections, then select nodes and move them to the target section. At least two nodes can form a group, and groups can be ungrouped; grouping is for organization and does not run automatically. Deleting a section puts its nodes back into the main section and keeps connections, media, and tasks. A separate free canvas is created from Manage canvases at the top, has its own saved versions, and its section number cannot replace the document number.

Find items using canvas content

Switch category, search by name / prompt / text, select a filtered result, and locate or edit it; the list and canvas point to the same object, not a second copy of the data. If you cannot find something, first check the section, category, and current document.

Save frequently used snippets as workflow templates

1. Select a set of regular text, reference, and generation nodes and save them as a template; internal connections and settings are saved with the template.

2. Set inputs such as {{主题}} for variable content. Publishing a new version keeps the old version and does not silently upgrade existing instances.

3. When inserting, fill in variables and reselect inputs of the same asset type; check the group, purpose, and references, then generate independently. An instance is not a live mirror of the template.

> Generation results and production business nodes are not brought into the template as copyable executable outputs. Inserting a template only creates an editable snippet; it does not generate or deduct model credits.

Local style and default model

Canvas settings can override the style and generation policy for this canvas; after clearing the override, it inherits the current work defaults again. The node chain button means follow defaults; manual settings apply only to that node. Work creation settings affect the shared baseline, so do not change the entire work for a single-shot experiment. Practice: set a local style in a test canvas, confirm the original canvas is not changed, then clear the override.

Canvas guide: models, specs, and real inputs

> Node generation settings → Model / Specs / References; video More settings → Input details

1. First select a provider and model allowed for the current workspace and work. The same model name may not have the same capabilities from the original provider and a third-party provider.

2. For images, choose an aspect ratio and size supported by the current model; for video, check the mode, duration, resolution, and combination limits, and recheck after switching

models.

3. Create references by uploading, dragging in, using @ tags, or explicit linking. Click a tag to preview it and verify order and purpose. Writing only a file name does not mean the file was passed in.

4. Before submitting, check the actual input list: which images, videos, and audio are sent to the model, and which are the first and last frames; invalid or unsupported references must be resolved. Do not assume the system will silently discard them.

Video mode	How to choose
Auto	With no references, it resolves to text-to-video; with references, to multimodal. The model must still support it; this is not a universal fallback.
Text-to-video	For text only; explicit text-to-video mode with references will be rejected.
Multimodal	Combine ordered images, videos, and audio according to model capabilities; distinguish appearance, motion, and camera-movement purposes.
First frame	Defines the starting image. The aspect ratio follows the first frame; this is not an ordinary composition reference.
First and last frames	Specify the start and end images separately; the model must support the corresponding count and mode.

Generation does not repeatedly show a cost confirmation; it checks workspace authorization, balance, and current platform rules. Text is billed too. The selected resolution is a request specification, not measured output pixels. For final shots, dialogue duration, continuation from the previous shot, and motion references still require dedicated checks. Exercise: use the same image to inspect a multimodal draft and a first-frame draft without submitting, and explain how their constraints differ.

Canvas guide: text candidates and shot drafts

Text node → Text purpose / history candidates / continue shot draft; Add → Shot draft

Save the text first, then turn it into a shot

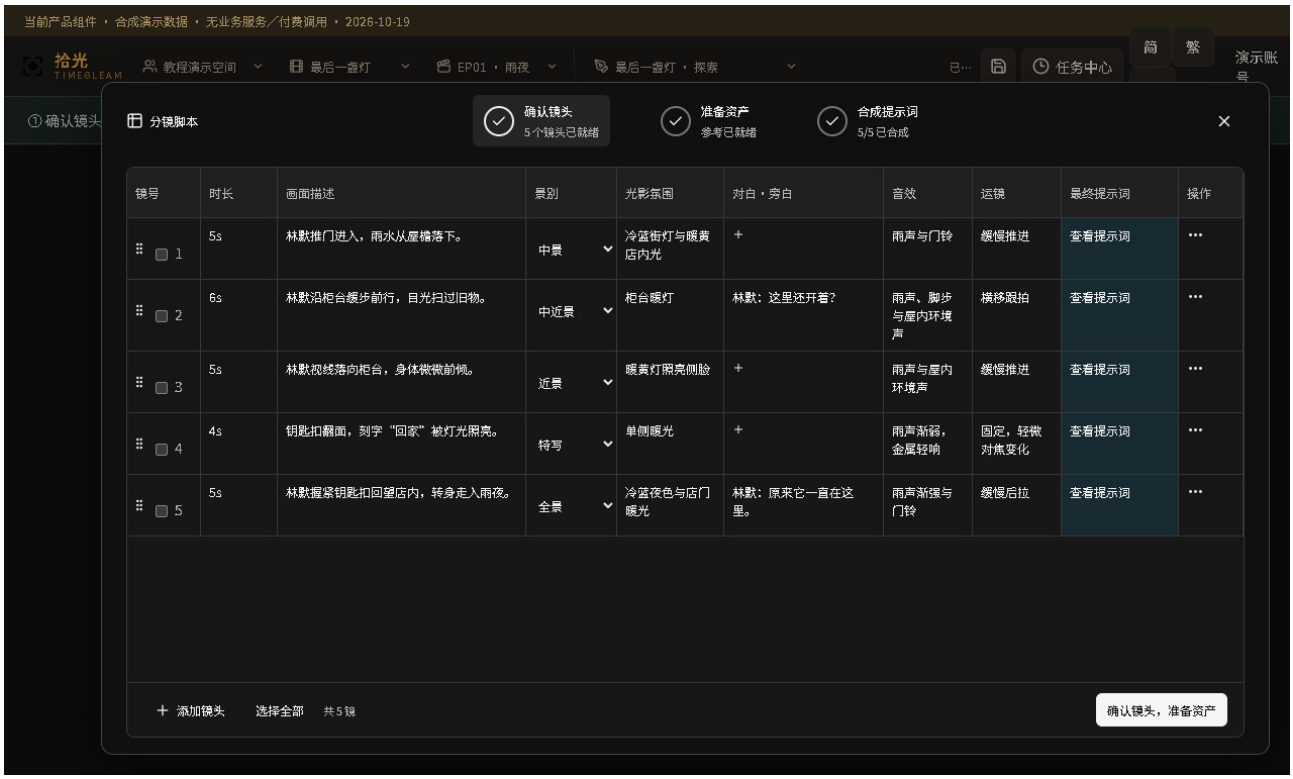
Text purpose offers general creation (default), simple rewrite, information extraction, and short script. The system selects a reasoning strategy based on purpose and integration capabilities, not by guessing from prompt keywords; for unverified providers, the model default is kept, which does not mean reasoning is disabled. After generation, browse timestamped text candidates, choose replace or append, and apply after checking the body text; paging through the preview does not change the body text.

Complete a free shot draft in three steps

1. Confirm shots: continue from the body text or add manually; edit the scene, characters / props, action, dialogue, shot size, camera movement, sound / music requirements, and suggested duration; add, delete, and move shots up or down.

2. Arrange references: select actual images for each shot, or explicitly check "No image reference for this shot". Source text snapshots are for tracing and are not silently rewritten when the source body text changes.

3. Prompts and generation: check the image and video prompt for each shot and confirm the current table; select the rows you need and generate the selected images or videos with the currently available model. Reuse successful and in-progress results; do not create duplicate tasks just to wait.



Confirm the content, organise references, then edit the prompt and generate selectively. Current components with example data, 2026-10-18; no model job was run.

Drafts are saved on this canvas and do not automatically become formal storyboards. For formal management, separately use "Adopt to workflow" and preview the selected rows. Web draft editable limits and external tool batch limits differ: up to 100 rows on the web, while a single MCP contract allows up to 20 rows; do not mix the two limits.

> Exercise output: keep the S01-S06 shot drafts, confirm the references and prompt for S04, and submit only that shot. Acceptance means checking the selected row, the real job and the resulting media, not that you clicked a button for the whole table.

Canvas guide: assets, subjects, and formal adoption

Asset center / Upload → media assets, subjects, characters / scenes / props; node result → Adopt to workflow

Media assets are files; a subject is a package of up to 12 images plus notes and audio; a formal visual asset is an identity with base/state versions. Subject references freeze the snapshot at the time; editing the package does not automatically refresh existing references. A formal version uses only one main reference image; auxiliary images and audio are not automatically written into formal character versions.

1. In Media assets, filter by type, category, and source; the Character / Scene / Prop tabs filter by identity, version, and base asset. The window can be dragged and resized, and reset to its default position; dragging an asset card creates a reference, not moves the window.

2. Upload and create a visual asset, or add references to an existing editable version. A state asset must be linked to a base version of the same type; locking the main image requires explicit confirmation. For a locked or deprecated version, unlock it or create a new version first; it cannot be written to directly.

3. Drag a regular reference image or a locked version into a video node to create an actual reference; if the same identity has a version conflict or a second scene, resolve it explicitly first - no silent replacement.

4. From free content, choose "Adopt to workflow". For first use, enable workflow and confirm the work name and style; choose a new or existing target, check the preview, then confirm adoption. Text can append/replace the script, drafts can select lines, and images and videos can become formal shots.

When upload succeeds but linking fails, retry linking with the uploaded material instead of uploading again. Adoption preserves the free source, old materials, and formal old versions; it does not auto-generate, replace old shot bindings, or change edit clips. For a standalone video, "Add to this episode's storyboard" also requires explicit confirmation.

Exercise: adopt the Lin Mo main image preview as a new visual version; check the impact first, then decide whether to confirm; then bind that version in the formal storyboard. A reference image appearing alone does not count as a completed formal binding.

Canvas handbook: all image tools

Image result → Preview / Edit → image toolbar; More tools; AI layer separation

First select an available source image. The workbench reads the source material and version and keeps the original file; save a copy after you are satisfied with the local preview. For model tools, read the current capabilities first and distinguish provider-native editing from reference-guided generation; the same name does not mean every model has native capabilities.

12 workbench tools and how each is completed

Tool	Actions, outputs, and cost boundaries
Smart image edit	Describe what you want to change and keep; draw a local mask when the model requires it; submit to generate new candidates and check whether the subject is distorted.
720 panorama	Enter the panorama creation path, select the corresponding model, and generate; a regular flat image does not become a panorama just by being renamed.
Outpaint	Adjust the target aspect ratio, expansion range, and description, check the frame preview, then generate; the original image is preserved.
Smart upscale	Select an available upscale factor and model capability, then check the actual pixels and detail after generation; it does not promise to restore information that is not there.
Cutout	Confirm the subject and supported models, and check the edges and transparency result; this is not a regular local crop.
Preset	Select a preset to create a draft, then confirm the parameters and generate; opening or inserting a preset alone does not call a model.
Multi-angle	Set the horizontal, pitch, and distance intent, and generate different viewpoints with reference guidance; this is not an exact reconstruction of the source image's 3D geometry.
Smart relight	Set direction, light color, intensity, and softness, and compare after generation; identity preservation and background still need to be checked.
Eraser	Draw a mask over the areas to remove, with brush strokes that can be added or removed; describe how to complete the surroundings, then check the occlusion boundaries after submitting to the model.
Crop	Locally adjust the crop box, rotation, and horizontal/vertical mirror; preview, then save as new without calling a model.
Grid split	Set rows and columns, select cells, and save the selected sub-images locally; first check whether the split lines cross the subject.

Tool	Actions, outputs, and cost boundaries
Annotate	Add strokes, lines, or boxes locally, adjust color and thickness, undo/redo, then save as new after confirming.

AI layer separation is a separate model-processing path that produces reusable layers; do not treat "mask painting," "grid split," and "semantic layer separation" as the same thing. All AI calls are billed by actual task; local processing adds no model fee. After the source image changes, recheck the workbench; do not force an old preview over the new image.

Exercise: first crop the keychain locally and save as new, then open outpainting to review capabilities and preview without submitting. The asset center should still contain both the original image and the new cropped image; downstream nodes must clearly specify which one they reference.

Canvas manual: voiceover, audio, and video sampling

Audio node generation settings; audio asset preview/edit; video asset reference handling; timeline soundtrack

A standalone audio node for voiceover is not a soundtrack generator

Enter the lines, choose an available voiceover service and voice, audition when a preview is available, then set speed and performance emotion. Without an audition, do not automatically start a billable generation. Locking a character's appearance does not mean the voice is ready; a subject's voice does not automatically become the official character voice. Create music in the timeline soundtrack; standalone sound effects can be uploaded.

Local audio tools

Choose a source, set trim in/out points, split, volume, fade in/out, and speed; preview first, then save as a new asset and keep the original sound. Analysis reads the loudness and silent intervals of the first audio track in the original file; it is not speech recognition and not multitrack mix analysis; it can be saved as a canvas analysis note.

Beat analysis provides onset/BPM candidates, and silence rough cut provides interval suggestions; rhythm editing also gives candidates first. Do not interpret candidates as already

automatically deleted clips or completed beat sync; review and apply them in the edit. Silenceis not necessarily useless content-a rainy-night pause may be the performance.

Video reference tools

Select still frames or clips by actual time and save them as references; check duration, in point, and source; keep the original video. Extracting a few frames does not mean a model has watched the entire video. For semantic shot analysis, use the next chapter's frame-by-frame review; ordinary sampling does not automatically call an understanding model.

> Exercise: trim an interval from your own rain audio, set a short fade in/out, and save as new; audition the original and the new asset separately. Generating voiceover or music is another explicit model call and is not billed together with local trimming.

Canvas manual: frame-by-frame review

> Add → Frame-by-frame review → Upload / select video from canvas / select from asset center



A detail of the current shot-analysis component with composited sources and demo authorisation, 2026-10-18. The 0 points in the image are a demo response from an isolated fixture, not a real quote, and do not mean model understanding is free.

1. Choose a video you have the right to analyze, up to 15 minutes and 50 MB. Play it to verify the actual source, and confirm there are no running tasks before replacing the source.

2. Select the shot, motion, and music dimensions and a currently authorized review model, then check the estimated cap. Unsupported dimensions block submission; do not assume every visual model can hear audio.

3. Explicitly choose "Start review" and check progress in the original task; interruptions can be resumed, and canceling does not reverse model usage already incurred. If some dimensions fail, retry only the failed dimensions and do not redo the successful parts.

4. In the source and analysis details, check each dimension, play the original video to verify, export this result, or use the grouped frame images, motion clips, and original-audio clips as references for later creation.

Local extraction is free; model understanding is billed by actual usage. Analysis is based on candidate sampling and does not mean every decoded frame is understood by the model. Music extraction preserves the original dialogue and sound effects and does not separate vocals/instrumentals; a successful task does not guarantee accurate music semantics.

The current default candidate is the licensed Seed 2.1 Lite (doubao-seed-2-1-lite-260915); it still needs to be enabled on the platform and authorized for the workspace. Qwen Omni is not a music candidate; it can be used for the text, storyboard, and motion it supports. If you select music, choose a model that supports audio understanding; it will not switch automatically on failure. After deleting a source or output node, task backfill cannot be used to revive the deleted node.

Exercise: choose a very short video you own and first check only the model, dimensions, and quote; after obtaining invocation authorization, you can analyze the storyboard separately. For acceptance, distinguish analysis results, extracted materials, and manually verified conclusions.

Canvas guide: Panorama and 3D director stage

Image panorama route / Add menu → 720 panorama, director composition; composition node → edit, continue, preview

A panorama saves the viewing direction

Choose a suitable panorama source, adjust orientation and framing in the viewer, and save the view/composition reference. Replacing the panorama asset is not the same as moving the camera; a regular image also cannot produce back-side content out of nowhere. Saving a composition does not mean video generation is complete.

Controllable elements in a director scene

1. Create a scene, place characters and props, and adjust position, size, orientation, and pose; when binding or unbinding a prop, verify the attachment target to avoid modifying other characters.

2. Adjust the camera, focal length/framing, and lighting, and save the still-frame composition. The assistant can also propose local changes; review the differences first without rebuilding the entire scene.

3. Set object/camera keyframes on the motion timeline, play the preview, and check paths, occlusion, the axis, and direction of travel. Automatic direction hints are only assistive and do not know the story's main axis.

4. Continue a new shot from a specified time and verify the source snapshot; after the source changes, use single-shot or batch "Update continuation shot" and choose item by item whether to keep this shot or apply the source difference. Only direct continuations propagate, not recursively across episodes.

Confirm stills, previews, and final bindings separately

Browser "Export motion preview MP4" requires the page to stay open; "Export motion preview in background" saves and freezes the scene first, and after the task is submitted you can close the page. The completed result serves as a silent motion reference. It does not call a video model or incur model fees, but it is not a final rendered film. If the environment or source changes, recover according to the actual task instructions; an old preview cannot be automatically backfilled into a new scene.

"Use as shot motion reference" requires selecting a real final shot from this episode and verifying the motion/camera-movement purpose; switching a non-multimodal shot requires explicit consent, and the original first/last images are kept as composition references rather than strict first/last constraints. Binding does not generate video, and unbinding does not automatically switch back to the old mode. Composition sync, continuation updates, previews, and final adoption each have their own impact confirmation and keep the old assets and edits.

Complex 3D, pose and multi-track motion editing is best done on desktop. Mobile screens can view and do common operations; this is not a promise that mobile has a full professional editor or has passed real-device acceptance.

Canvas guide: candidates, tasks, and three histories

Node assets and versions; top generation history; task center; bottom save history

Entry point	How to look up and use
Node candidates	View successful results from the same generation node and select two for side-by-side comparison; previewing is not adoption. After confirming the impact, adopt the new candidate; old media is kept, downstream items can be marked as needing update, and regeneration is not automatic.
Generation history	View the latest 100 by canvas/episode/work, type, and status; locate the source; read frozen inputs and text results; reuse assets or create a variant draft. A variant is not a resubmitted task.
Task center / Run history	View queued, running, failed, pending recovery, and canceled statuses; check task ID, cause, and cost receipt. Closing the page does not delete background tasks.
Save history	Recovery restores a document content version, saved as a new revision; it does not roll back provider charges or delete later files. It is entirely different from generated candidates.

On failure, check the original task first; an unknown acceptance status must not be treated as a definite failure and then used to open another task. When recovery is possible, reuse the original task and remote receipt; when billing authorization is needed again, read the prompt carefully first. Canceling a task already submitted to the provider does not guarantee billing stops. After adopting a new video, old timeline instances still point to old assets; verify and update them separately on the timeline.

Exercise: compare two candidates on the same node and note the currently adopted asset ID; preview only without confirming-the current adoption should stay unchanged. After creating a variant draft, check the task list; no generation task should appear out of nowhere.

Canvas guide: Assistant and batch production

Left AI assistant / Ctrl+/; More tools → Batch generate, Batch production, Run history

State the scope clearly

Check the conversation title and scope before choosing references: up to 8 fixed nodes, with follow selection optionally enabled; remove invalid or expired references before sending. Conversations can be searched by name, created, renamed, and switched; activity times are shown based on server records, and long historical answers can be collapsed. Conversations in different free documents are isolated; do not expect the assistant to automatically know about another canvas.

"Arrange a creative draft / split into shot drafts / adjust director scenes / draft a generation plan" only fills the draft; you need to send it yourself. The assistant can give plans and editable proposals and diffs; review, adopt, and save them, then explicitly generate. An answer saying " suggested completion" is not a save receipt, let alone a model output.

Batch generation and production plans

First determine the selected nodes, real upstream, and reused items. Batch generation submits the selected tasks; batch production / running a workflow advances dependencies according to the plan-verify whether existing content may be modified, the selected model, and the cost range. Saving a production plan is not the same as starting it. During execution, view steps in run history; on failure, recover or handle the original step first. Pausing prevents further advancement but does not revoke in-flight calls. Do not bypass conflict or duplicate-charge protection by rebuilding the plan.

Voice and text share the same execution boundary

Clicking voice call explicitly is required to connect the microphone, and a secure origin and authorization are required; voice input transcribed into the input box still needs to be sent. Collapsing the panel or stopping an answer does not cancel submitted tasks. Text, voice, and background plans are all subject to this round's tools, permissions, references, version, and impact confirmation.

Check only the S04 I reference and its direct inputs: list missing references, the current model mode and the reason an update is needed. Give a repair proposal without executing it first; do not generate, do not switch models and do not touch other shots. When a change is needed, let me check the differences before saving.

Canvas guide: One-click film pricing and independent editing

Home canvas creation → One-click canvas plan / One-click canvas film; original canvas → Production pricing and film nodes

Choose specs first, then pay the planning fee

1. Enter creative requirements; add up to 12 references as needed. "One-click canvas plan" only plans and creates editable nodes and connections; "One-click canvas film" plans first, then calculates subsequent production. Planning is also charged; it is not a free trial.

2. Choose a model that is authorized, enabled and suited to the purpose in the current workspace; an expired old configuration cannot be treated as a valid default. In the main area, check the video model, video generation resolution, aspect ratio and target total duration; " More settings" contains the text/image/voiceover models for the plan, single-shot generation duration, reference method and export resolution. Specifications vary by model; when " Confirm current valid specifications" appears, review them before confirming.

3. The target total duration is for planning up to 8 shots, not the model's per-shot duration. Exporting at 720p/1080p does not increase native generation clarity; if reference materials are not applicable, explain the reason explicitly and do not remove them silently. Soundtrack uses uploaded audio; voiceover is generated only when needed and the model is available-music is not generated automatically.

4. First review the planning token unit price, the fallback price when there is no valid usage receipt, and the hold details, then click "Generate plan and price" (the plan entry point is " Generate canvas plan"). Planning is held first, then settled by actual usage and the difference is released; before shot tasks are finalized, an exact price for the whole film cannot be promised. The plan is saved as editable text, shot drafts, media/subjects, connections and groups.



Check the model and specifications first, then run the charged planning step. 2026-10-19, real components with example settings; the model names and prices in the image are example data and do not represent real authorisation or pricing.

Subsequent costs require separate confirmation

1. After planning is complete, check "Planning settled", the number of shots, pending generations, reused materials and "Cost details". Settled planning points are listed separately and are not counted again in later holds; existing successful materials are reused first. Local editing and export add no model fee.

2. After confirming "This run holds ×× points", click "Start film"; only the tasks approved this time are authorized. The home page no longer requires a manually entered point cap for the whole film; the system dispatches based on the scope confirmed this time and places holds item by item, releasing failed items according to the settlement result-this does not mean the full quote is charged immediately or that you authorize spending your entire balance.

3. Quotes are valid for 15 minutes. After changing the model, specifications, canvas content, references or price, you must "Re-price"; if canvas content changes, first "Update affected content". If the balance is insufficient, references are invalid, or the original call still has a hold, follow the prompt, then re-price-do not keep creating new plans.

4. Only the original initiator can confirm production costs. If the plan is already saved, continue from the home page "Continue saved canvas plan · Skip planning", the dialog "Enter canvas to view plan", or the original canvas "Production pricing" without planning again. Successfully generating only the plan does not mean the subsequent media has been produced.



The example separates plan settlement from the later hold; this production is dispatched only after confirmation. 2026-10-19, composited shots, points and quote, not the measured cost of The Last Lamp; the preview page forbids submitting a confirmation.

Pause, edits and old records

The progress bar shows completed steps, settled items, in-flight holds and the amount authorized this time; "..." shows the model, references, steps, failure location, locked materials and export history. Authorized production dispatches downstream only after the upstream actually succeeds; closing the browser does not cancel authorized production, pausing stops only further dispatch, and submitted tasks continue to settle. Other editors can pause/cancel but cannot continue cost authorization on behalf of the original initiator.

After edits, "Update affected content" keeps unrelated successful materials and locked results, then re-price and confirm; moving, scaling and color do not trigger generation. If a request's acceptance is unknown, first restore the original task and check the receipt; do not automatically resend paid requests. When a fixed model configuration changes, stop dispatch; do not silently switch models or drop references. Old records still keep the original point cap and the "Save limit and pause" rule; do not treat old-record controls as required steps on the new home page.



Open the original production record instead of creating another work; the progress bar's budget comes from this authorisation, and a new record continues through Make-film quote. 2026-10-19, composited status and points, with no real jobs, provider results or delivered film.

Film projects are independent of the official episode timeline

The film card can play, download, and open a standalone edit, where you can adjust shot order and trims, subtitles, and audio mixing; save and re-export. Changing volume and subtitles does not regenerate the video. Manual edits are not rebuilt automatically; "Rebuild rough cut" replaces the film's manual edits, requires explicit confirmation, and keeps historical exports. When exiting with unsaved edits, save first or explicitly discard. Each canvas allows up to 20 production records; each record allows up to 20 historical exports.

> One-click film does not adopt into the formal workflow automatically and does not rewrite the formal episode timeline. Only a successful export whose file plays counts as a finished film; a queued job ID or a demo screenshot is not real generation acceptance. If you only want a draft, choose New blank canvas or a manual entry under More creation, not the automatic planning or film buttons.

Canvas guide: editing, subtitles, checks, and delivery

Top timeline / Add result to timeline → editor; film node → Open standalone edit

1. Add assets that are explicitly applied or selected; arrange video, images, text, subtitles, and audio by track. Confirm whether you are in the main timeline or a standalone film project, and do not mix them.

2. Adjust position, trim, in point, speed, frame layout, volume, and fades per clip; use keyframes/envelopes, transitions, effects, and color grading according to the actual panel. Preview uniform transitions and background music before applying, so they do not mask dialogue.

3. For music generation, choose the currently available model and parameters; after explicit submission, charges follow the model rules. Regular stitching, volume adjustment, and local export do not add model fees. AI video processing such as Topaz Proteus is checked against the corresponding task and is not equivalent to free editing.

4. Create subtitles and import/export SRT or WebVTT as needed; a subtitle file does not prove the subtitles are burned into the video. Preview the full result, then choose the export you need.

5. After a shot adopts a new version, click "Check for new canvas versions", select the clip instances that need updating, and preview the impact before confirming. Preserve timing and compatible effects where possible; unselected instances, old assets, and old exports are not replaced automatically.

6. Run "Edit health check", resolve current check items such as missing sources and out-of-range segments; preview the beginning, middle, end, and audio, then export and play the downloaded file.

Project packages are for continuing work, not the final video

Export/restore project packages can move an edit project; before restoring, verify the source and compatible assets, and do not overwrite unsaved edits. When another browser or MCP saves a new timeline, save the local draft first or handle the conflict; do not silently overwrite from an old page. Export is an independent task after saving a snapshot; closing the page, it continues according to task status. If it fails, check the original task. After download, verify duration, aspect ratio, audio/video, and the file itself; keep old exports for comparison.

Exercise: add S04 in two places, update the candidate and replace only the first instance; keep the second on the old version. Save the subtitles and the project package separately, restore them on a copy, and finally download one clearly versioned video.

Canvas guide: saving, permissions, mobile, and cleanup

Top save status; bottom save history; More → Trash; Asset center → Clean up assets

For conflicts, preserve both sides' content first

Normal autosave and manual save are both protected by server-side revisions. When the remote has changed, autosave pauses; read the changes first and merge compatible fields. For same-field conflicts, keep the local draft or download a backup, then choose explicitly. Restoring history is saved as a new version; do not refresh and hope. Unsaved input on the web page cannot be read by MCP; save first before continuing across devices.

The three deletion options have different consequences

action	Retention and risk
Delete node / remove reference	Original media is not deleted; regular objects can be undone on the canvas.
Recycle canvas / episode / work	Soft-recycled per page; history and references remain. Recycling does not mean disk space is freed.
Permanent cleanup	Workspace owners/admins can permanently delete recycled works, production sets, or free canvases, or clean up assets with no references; regular editors and read-only members cannot.

Permanently deleting an entire work includes all episodes, free canvases, recycled subsets, historical versions, and dedicated records; you can also permanently delete a recycled episode or a single free canvas. An episode includes its workflow, main canvas, and associated free canvases; the main canvas cannot be permanently deleted separately. Cleaning up a single free canvas keeps the episode workflow and other canvases. If the parent has already been recycled, handle the parent scope first; shared assets or assets that still have references are retained or blocked from deletion based on the impact preview. If no valid production set is

available to host it, retain or restore one episode first. There is no self-service permanentdeletion for accounts; for permanent workspace deletion, see the owner management notesbelow. Asset cleanup checks current, historical, recycled, adoption records, and unfinishedtasks; if references exist, handle them first, and it does not automatically unlink andforce-delete. The asset center is limited to the current work and lets you check names,references, and batch selection; the web preview cannot undo the impact. After one explicitconfirmation, processing runs in the background; closing the window does not cancel it, andfailures can be retried in Permanent cleanup records.

> Permanent cleanup cannot be restored from the trash and does not refund settled generation fees; necessary accounting records are kept, shared files stay while other materials use them, and capacity release is not promised. The assistant and MCP cannot bypass the web impact confirmation to run new permanent cleanup. The exercise only checks the preview and blocking reasons; do not use real deletion as acceptance.

Continue on mobile; refine professionally on desktop

The mobile entry keeps the space, work, stage and canvas; modes use a dropdown at the top and the bottom assistant and production bar handle common tasks. Touch works through node details and does not require right-click or desktop shortcuts. Complex director stage, image retouching and multi-track production are best done on desktop; this tutorial's narrow-screen checks do not represent real iOS/Android devices or acceptance of every professional tool.

Workspace management, point refunds, and entire workspace deletion

Top workspace dropdown → Manage workspaces: view the workspace list first, then expand to enter workspace info, members and invitations, workspace points, tasks and queues, activity log, recycle bin and cleanup, and dangerous actions. Viewing a workspace other than the current one does not switch your creative workspace; to switch, click "Switch to this workspace." "Me" keeps personal profile, personal points, invitations, and account security; the allocation shortcut in workspace points goes to the personal allocation form.

Only an active workspace owner can return points that the original contributor has not consumed, reserved, or let expire and that can be traced back to their personal wallet. Enter a single contributor, amount, and reason, preview the recipient ID, amount, and remaining workspace balance, then confirm; the original validity period is unchanged, no cash is refunded, and no tasks are canceled. Users who have left the workspace but still have a valid account can still receive points; abnormal accounts or sources require platform verification. Administrators' existing permission to clean up recycle bin contents is unchanged.

Whole-space permanent deletion exists only on the owner's danger page: preview every member, work, production set and main canvas, free canvas, material, history, trashed item, job and expected point refund, then type the full space name once to confirm. The system freezes new operations, waits for settlement, refunds by source and cleans up, and cannot be undone; necessary accounting, audit, backups and shared files still referenced are kept, and capacity release is not promised. Deleting the last space does not delete the personal account; you can still open the personal wallet, invitations, security, space creation and the space-deletion receipts you started. Accepted is not completed; this section's exercise checks the preview only - never use a real business deletion for acceptance.

08 More Tools and Troubleshooting
Workflows, adaptation, extension tools and common questions

Workflow guide: final scripts and batch production

Multi-episode storytelling, final assets, and batch management follow this path. The old nine-lesson links remain available below; old progress is kept and does not count toward the new eight lessons.

Full flow and views

Script / production prep → visual assets and storyboard planning → shot production (composition optional) → editing → delivery. The workflow canvas is another view of the same episode's production data.

Create a work and configure it

Home → Workflow mode → New work, then choose direct creation or story adaptation. Existing free canvases can explicitly enable workflow; the first time, add the work name and style. The creative bible is shared; this episode's supplementary requirements are saved per episode; click Confirm after editing.

Script and production prep

Direct creation can be saved manually, imported intelligently, or generated/revised by AI. Story adaptations are generated from the STORY page; the Script Room can still save manually. Have a valid final script before storyboarding.

In Production Prep at the bottom, select this EP, text, and image models, then submit " One-click breakdown and generate storyboards and assets." Switching the current episode

does not change the selected scope. It handles breakdown, storyboards, assets, bindings, andapplicable voices; it does not generate video. "System adopted · Pending review" still requiresmanual image review.

> The Last Lamp
> Rainy night. An old grocery store on the corner is about to close.
> Lin Mo pushes the door open and enters, folding his dripping black umbrella. The counter light is dim yellow.
> He sees an old brass keychain under the counter glass and stops. "Home" is engraved on the back.
> He picks up the keychain, presses the money onto the counter, and turns off the last lamp at the counter.
> He walks out of the store, holding the keychain, and goes down the wet street into the night.

Go to script

Breakdown and storyboards

In the manual route, complete the script breakdown first; valid source results are adopted automatically, and you no longer need to click to confirm the breakdown. A standard breakdown only creates requirements; it does not generate or lock images. After Production Prep is complete, check the results directly; do not rerun it.

In the bottom action area, select the episode and model. Blank episodes automatically adopt the latest valid plan; for existing storyboards, compare first, then explicitly adopt while preserving media. Check characters, scenes, actions, shot size, camera movement, audio, and duration.

Breakdown requirements · Storyboard planning

Visual assets and character voices

State references use the correct base version. Manual generation follows the current list and filters, and existing images are skipped; state images must lock the parent reference. After previewing, confirm and lock the current list. Upload and create assets, or choose reusable images from media.

When there is dialogue, "Fill in character voices" and preview first. Unlocking may create a derived version; old bindings do not upgrade automatically. Subject association does not equal final adoption.

Visual assets

Optional compositions and storyboard images

New shots can generate video directly from bound visual assets. When needed, review the composition before shooting video; strict first frame / first and last frames require images. Verify identity, costume, space, and subject position.

Compositions and storyboard images

Final shots and candidates

Review the final submission and reference list, and handle missing images, unlocked items, and invalid bindings. Select the scope to generate video and explicitly adopt candidates; pending update does not mean reshot.

"Adopt to workflow..." requires previewing the target, identity, and version before confirming. The final subject version adopts only the main image; auxiliary images and audio stay in the asset pack. For a video to become a final shot, also confirm the target. Adding to the canvas timeline does not equal final adoption.

Shot production

Editing, review, and delivery

Place final shots into editing, and adjust order, trims, transitions, and audio. After applying " Unified transitions / background music," preview in full, then export. On the Delivery page, start checking this version's file for decoding, duration, black frames, still frames, and audio tracks; recheck after re-exporting.

AI frame spot checks and Topaz Proteus enhancement are optional and may incur charges; they require an available service. Save a copy and keep the original. Checks do not replace a full manual viewing.

Editing · Delivery

Assistant and execution status

The text assistant and live conversation share the same actions, subject to permissions, references, and confirmations. To modify a final script segment, first select the fields to adopt; visual asset drafts are saved after review and do not equal generating reference images. Saving a production plan does not start it; background execution begins only after explicit authorization.

Stopping a response, closing the assistant, or interrupting voice does not undo saved changes or submitted tasks; for pause and resume, check the corresponding task / plan.

Start production from STORY

Novels and synopses start from STORY materials. Adding, importing, and saving do not call the model automatically, and there is no need to fill in the old adaptation plan. Select continuously from the first unfinished chapter; total episodes are the combined episode count generated from all chapters, and duration is entered in seconds. The model analyzes episodes in the background; later chapters open only if the whole batch succeeds, and on failure, restore the original task first. Chapters and EP do not need to correspond one-to-one; after the script is complete, return to production prep.

Open STORY materials

More creation tools

Find tool entry points and usage boundaries here; to learn shot, narrative, and review methods, see Advanced creation, and keep the two separate.

Continuous shots and partial redo

A final shot using "Continue from previous shot" uses the previous shot's valid video and keeps this shot's assets; it is unavailable for the first shot and cannot use independent motion/camera references at the same time. Upstream changes make downstream items pending update but do not delete old videos. Whether a regular connection can reach the model depends on supported items; this does not guarantee a continuous result.

Panorama, 3D composition, and asset editing

The composition tools in the Add menu provide panorama framing and 3D composition. For video, extract frames and clips; for images, crop and annotate; for audio, adjust the selection, volume, speed, and fade in/out. Preview before saving as a new copy, keep the original asset, and use a computer for complex operations.

Batch production and background plans

Free canvas More tools → Batch production runs according to saved nodes and dependencies. Target-only and redo-upstream ranges are different. For checkpoints, retry after failure, and pause/resume, see the original run record; task success does not mean it has been reviewed or adopted.

Canvas management and color location

Manage canvas lets you rename, duplicate, and recycle free canvases. A color mark selects and centers elements of the same type without making results stale; read-only lets you locate but not modify. Dragging the edit card title moves only the panel; adjust node positions separately.

Open workflow canvas

If you hit a problem, continue here

Locate first, then decide whether to regenerate

Symptom	Check first
Character changes in every shot	Main image, actual references, final version, and binding.
Motion or composition is wrong	Whether there are too many requirements, the reference method; change only one thing at a time.
Preview is new, edit is old	Current adopted item and clip source material; clarify the replacement scope.
Failed / pending recovery	Original task cause and recovery entry point; reuse successful results.
No permission / insufficient balance	Account, workspace, role, credits; models are managed by platform admins.
Old content after refresh / conflict	Keep the draft, check the remote, and use the existing merge or backup entry point.

Mobile instructions

The mobile home page can quickly make a short film or create a regular work. The bottom production bar edits text, redoes one shot, adopts candidates and exports; the mobile assistant describes requirements. Edit this shot with the assistant brings in references but you still send it yourself. The free canvas uses touch and node details; desktop double-click or right-click is not required on mobile. Narrow screens collapse modes into a dropdown; the 3D director stage and dialogue/multi-track finishing continue on desktop.

Voice calls and permissions

In the current session, the microphone connects only when you click voice call; opening the assistant alone does not start the microphone. It requires HTTPS or localhost, microphone permission, and recognition/synthesis services. Text and voice share history, references, and receipts; after voice input goes into the input box, you still need to send it manually. An explicit goodbye is a local hang-up command, and collapsing the call keeps a visible control bar. Conversation history and unsent drafts are saved by user, episode, mode, and session; check references before switching.

Learning records and versions

The new eight lessons are saved only in this browser and do not sync across accounts or devices. The previous nine lessons are kept as-is, with their count shown; new lessons are not marked complete automatically; resetting this version does not delete old records. If storage is disabled, a notice appears, but you can still read and copy.

Version b3a7988c, date 2026-10-19, covering the unified episode canvas and creation tools. Images use current components with example data, not real paid generation. Feature availability depends on the deployed service version. The PDF and the web edition share one source; release 0.6 is kept only for older-version reference.

09 Tool Index
The full remote tool contract

MCP: complete index of 85 tools

The following is generated from the current registry and lists purpose, input fields, and practical positioning for each item. For field constraints, nested types, and optional defaults, use MCP inputSchema, tools.schema.json, or offline tools.schema.json as the complete reference. A read-only marker means it does not change business state and does not replace server-side scope checks; all writes, generation, and downloads remain subject to permissions and the security boundaries in this chapter.

identity_get

Verify the workspace bound to the credential, the current role, and service capabilities; do not display the token.

Required: none.Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · identity_get Returns the corresponding walkthrough.

workspace_list

List the workspaces visible under account-wide consent, with stable IDs and effective member permissions; consent covers future memberships.

Required: none. Optional: none.

Read-only/preview operation. Exact contract: MCP inputSchema · workspace_list； Returns the corresponding walkthrough.

workspace_switch

Select an authorized workspace for this conversation; business calls then carry its workspace_id each time and leave web or other chats unchanged.

Required: workspace_id (string).Optional: none.

Read-only/preview operation. Exact contract: MCP inputSchema · workspace_switch； Returns the corresponding walkthrough.

server_capabilities

Verify the service release identifier, actual tool schema, contract fingerprint, and current connection permissions to avoid using stale caches by mistake.

Required: none.Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · server_capabilities; Returns the corresponding walkthrough.

model_list

Read the providers and model IDs allowed in this workspace; does not return keys.

Required: none.Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · model_list; Returns the corresponding walkthrough.

production_list

List visible works; locate the real work ID first.

Required: none.Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · production_list; Returns the corresponding walkthrough.

production_get

Read the work's shared context, characters, and scenes to avoid creating another identity with the same name.

Required: production_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · production_get; Returns the corresponding walkthrough.

episode_list

List a work's episodes and their real numbers; do not build IDs from display episode numbers.

Required: production_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · episode_list; Returns the corresponding walkthrough.

episode_get

Read the saved document, scope, nodes, references, timeline, and dual revisions.

Required: project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · episode_get; Returns the corresponding walkthrough.

project_list

List production documents accessible in the credential's workspace; continue using scope to verify ownership.

Required: none.Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · project_list; Returns the corresponding walkthrough.

project_create

Create a work and its first episode; name and style must be explicit; retrying with the same request_id does not create duplicates.

Required: request_id (string)、name (string)、style (string).Optional: creation_mode、brief、video_ratio、video_resolution、video_reference_mode、duration、workspace_id.

Writes state, submits tasks, or saves local files; verify authorization first. Exact contract: MCP inputSchema · project_create; Returns the corresponding walkthrough.

episode_create

Create an episode in the specified work while keeping other episodes; if the response is lost, check the original receipt first.

Required: production_id (string)、request_id (string)、title (string).Optional: workspace_id.

Writes state, submits tasks, or saves local files; verify authorization first. Exact contract: MCP inputSchema · episode_create; Returns the corresponding walkthrough.

script_list

Read the work's current official script; free-text drafts are not treated as the official script.

Required: production_id (string).Optional: workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · script_list; Returns the corresponding walkthrough.

asset_list

Find media by scope, kind, view, episode and section; candidate files are still available here.

Required: project_id (string).Optional: scope、kind、view、episode_id、canvas_id、workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · asset_list; Returns the corresponding walkthrough.

receipt_get

Look up receipts for successful creates, uploads or plain commands in your own workspace; a 404 does not mean you can switch to a new ID.

Required: request_id (string).Optional: workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · receipt_get; Returns the corresponding walkthrough.

creation_guide

Read the new-video, retake or resume guide offline; nothing is produced.

Required: workflow (string). Optional: none.

Read-only / preview operation. Exact contract: MCP inputSchema · creation_guide; Returns the corresponding walkthrough.

shots_create

With both revisions, add 1-20 official shots at once, keeping existing shots; this is not a whole-table replace.

Required: project_id (string)、revision (integer)、production_revision (integer)、shots (array)、request_id (string).Optional: canvas_id、workspace_id.

This writes state, submits jobs or saves local files; check authorization first. Exact contract: MCP inputSchema · shots_create; Returns the corresponding walkthrough.

shot_update

Change only one official shot and its image/video settings; parameters are merged and previous results are kept.

Required: project_id (string)、shot_id (string)、revision (integer)、production_revision (integer)、request_id (string).Optional: patch、image、video、workspace_id.

This writes state, submits jobs or saves local files; check authorization first. Exact contract: MCP inputSchema · shot_update; Returns the corresponding walkthrough.

references_set

Maintain images uploaded through this tool as references for one generation node, up to 8; other relations are kept.

Required: project_id (string)、node_id (string)、revision (integer)、production_revision (integer)、asset_ids (array)、request_id (string).Optional: workspace_id.

Writes state, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · references_set Returns the corresponding walkthrough.

generation_preview

Review the actual generation plan for the selected 1-20 nodes, without submitting or reserving.

Required: project_id (string)、node_ids (array).Optional: exact、workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · generation_preview Returns the corresponding walkthrough.

job_list

List the 200 most recent jobs in the current document; use the persistent job_id to look up older jobs.

Required: project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · job_list Returns the corresponding walkthrough.

job_get

Read the original job's status, progress, results, errors, and cost; when unknown, check here first.

Required: job_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · job_get Returns the corresponding walkthrough.

job_cancel

Request cancellation; usage already submitted to the provider is not guaranteed to stop or be refunded.

Required: job_id (string).Optional: workspace_id.

Writes state, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · job_cancel Returns the corresponding walkthrough.

candidate_list

List the successful image/video candidates retained for the node; they are not adopted automatically.

Required: node_id (string)、project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · candidate_list Returns the corresponding walkthrough.

candidate_preview

Preview adoption and downstream impact using the real nodeId, jobId, and assetId.

Required: project_id (string)、nodeId (string)、jobId (string)、assetId (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · candidate_preview Returns the corresponding walkthrough.

candidate_adopt

Add the original candidate selection with an explicit impact_token to adopt it, keep the media, and do not replace the edit instance.

Required: project_id (string)、nodeId (string)、jobId (string)、assetId (string)、impact_token (string)、request_id (string).Optional: workspace_id.

Writes state, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · candidate_adopt Returns the corresponding walkthrough.

timeline_get

Read the saved timeline and revisions; first distinguish them from unsaved web drafts.

Required: project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · timeline_get Returns the corresponding walkthrough.

timeline_preview

Preview basic edits such as adding tracks/clips, trimming, speed, and text, while preserving unrelated tracks.

Required: project_id (string)、request_id (string)、command (object).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · timeline_preview Returns the corresponding walkthrough.

timeline_apply

The original request_id and command are added to the preview credential on apply; replaying the original request returns the current document, and a new request using an old preview causes a conflict.

Required: project_id (string)、request_id (string)、command (object)、impact_token (string).Optional: workspace_id.

Writes state, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · timeline_apply Returns the corresponding walkthrough.

timeline_media_scan

Finds timeline instances that still reference previously adopted media; scanning does not replace them.

Required: project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · timeline_media_scan Returns the corresponding walkthrough.

timeline_media_preview

Select elementId and nodeId; the preview shows only the impact of updating those instances.

Required: project_id (string)、request_id (string)、selections (array).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · timeline_media_preview Returns the corresponding walkthrough.

timeline_media_apply

Updates the reviewed instances, preserving timing and compatible effects; other instances remain unchanged.

Required: project_id (string)、request_id (string)、selections (array)、impact_token (string).Optional: workspace_id.

Writes state, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · timeline_media_apply Returns the corresponding walkthrough.

export_preview

Validates the current saved edit, duration, specs, and media, and freezes the export intent.

Required: project_id (string)、request_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · export_preview Returns the corresponding walkthrough.

export_submit

Submits a previewed local composite and returns a persistent job; the same snapshot and request are reused without calling the model.

Required: project_id (string)、request_id (string)、impact_token (string).Optional: workspace_id.

Writes state, submits a task, or saves local files; verify authorization first. Exact contract: MCP inputSchema · export_submit Returns the corresponding walkthrough.

creation_draft_create

Creates a draft work, first episode, and free canvas without generating; intent can be text/blank.

Required: request_id (string).Optional: intent、brief、name、workspace_id.

Writes state, submits a task, or saves local files; verify authorization first. Exact contract: MCP inputSchema · creation_draft_create Returns the corresponding walkthrough.

canvas_list

Lists same-episode workflow views, standalone free documents, and workflow setup status.

Required: project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · canvas_list Returns the corresponding walkthrough.

canvas_create

Creates a free document with a production episode ID; source_id can copy the same episode's canvas and strip the task execution identity.

Required: project_id (string)、request_id (string).Optional: name、source_id、workspace_id.

Writes state, submits a task, or saves local files; verify authorization first. Exact contract: MCP inputSchema · canvas_create Returns the corresponding walkthrough.

workflow_setup_get

Reads the name, style, and current revision required by the production workflow; verify before enabling.

Required: project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · workflow_setup_get Returns the corresponding walkthrough.

workflow_enable

Explicitly enables the production episode workflow it belongs to; does not generate or adopt free content.

Required: project_id (string)、revision (integer)、production_revision (integer)、name (string)、style (string)、request_id (string).Optional: workspace_id.

Writes state, submits a task, or saves local files; verify authorization first. Exact contract: MCP inputSchema · workflow_enable Returns the corresponding walkthrough.

subject_list

Reads this work's asset package, source documents/nodes, and saved versions; does not bind or adopt.

Required: project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · subject_list Returns the corresponding walkthrough.

canvas_command_preview

Previews ten partial commands for ordinary objects and returns creation IDs and impact; does not accept arbitrary whole-document writes.

Required: project_id (string)、request_id (string)、revision (integer)、production_revision (integer)、command (object).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · canvas_command_preview Returns the corresponding walkthrough.

canvas_command_apply

Applies with both revisions, the original request_id and command plus impact_token; replaying the same request makes no duplicate changes.

Required: project_id (string)、request_id (string)、revision (integer)、production_revision (integer)、command (object)、impact_token (string).Optional: workspace_id.

Writes state, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · canvas_command_apply Returns the corresponding walkthrough.

canvas_recycle

Soft-recycles a free canvas using the latest dual revisions, preserves media and formal content, and a retry returning not_found does not restore it.

Required: project_id (string)、revision (integer)、production_revision (integer).Optional: workspace_id.

Writes state, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · canvas_recycle Returns the corresponding walkthrough.

video_specifications_get

Reads the server's current model mode, duration, resolution, and limits; do not guess specs.

Required: project_id (string)、provider_id (string)、model (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · video_specifications_get Returns the corresponding walkthrough.

video_reference_capabilities_get

Reads the model's capabilities for motion and multimodal reference combinations, counts, and durations.

Required: project_id (string)、provider_id (string)、model (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · video_reference_capabilities_get Returns the corresponding walkthrough.

video_input_preview

Validates the ordered media list for a saved video node with the same actual compiler; invalid references are explicitly rejected.

Required: node_id (string)、project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · video_input_preview Returns the corresponding walkthrough.

adoption_targets_get

Lists valid formal targets in the episode that owns the source free canvas; cross-episode adoption is not permitted.

Required: project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · adoption_targets_get Returns the corresponding walkthrough.

adoption_preview

Preview visual/shot/storyboard/script adoption; freezes the source, target, and media versions.

Required: project_id (string)、request_id (string)、source_node_id (string)、kind (string).Optional: card_id、visual_kind、parent_version_id、name、description、lock_reference、shot_id、row_ids、script_mode、workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · adoption_preview Returns the corresponding walkthrough.

adoption_apply

Confirm formal adoption with the original selection plus impact_token; keeps the source's previous version, does not generate or change the edit.

Required: project_id (string)、request_id (string)、source_node_id (string)、kind (string)、impact_token (string).Optional: card_id、visual_kind、parent_version_id、name、description、lock_reference、shot_id、row_ids、script_mode、workspace_id.

Writes state, submits tasks, or saves local files; verify authorization first. Exact contract: MCP inputSchema · adoption_apply Returns the corresponding walkthrough.

adoption_receipt_get

Look up the formal adoption receipt with the original source document and request_id to avoid adopting twice after a disconnect.

Required: request_id (string)、project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · adoption_receipt_get Returns the corresponding walkthrough.

work_search

Paginated search for works in the workspace bound to the credential; does not access assets in other workspaces.

Required: none.Optional: q、offset、limit、workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · work_search Returns the corresponding walkthrough.

work_entries_get

Reads the real entry IDs for a work's episodes and canvas.

Required: production_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · work_entries_get Returns the corresponding walkthrough.

film_list

Lists the films already saved on the current canvas; does not start planning.

Required: project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · film_list Returns the corresponding walkthrough.

film_get

Reads the specified film, revisions, shot bindings, affected steps, and past exports.

Required: film_id (string)、project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · film_get Returns the corresponding walkthrough.

film_options_get

Read the models available in this workspace and the current production specs.

Required: project_id (string).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · film_options_get Returns the corresponding walkthrough.

film_create

Create a film by reusing a complete plan or existing storyboards, without replanning or authorizing generation.

Required: project_id (string)、request_id (string)、revision (integer)、production_revision (integer)、brief (string).Optional: plan、draft_node_ids、asset_ids、targets、ratio、duration、video_resolution、shot_duration、reference_mode、export_resolution、workspace_id.

Writes state, submits tasks, or saves local files; check authorization first. Exact contract: MCP inputSchema · film_create Returns the corresponding walkthrough.

film_manual_create

Create a standalone film edit from existing assets without calling a model.

Required: project_id (string)、request_id (string)、revision (integer)、production_revision (integer).Optional: name、asset_ids、workspace_id.

Writes state, submits tasks, or saves local files; check authorization first. Exact contract: MCP inputSchema · film_manual_create Returns the corresponding walkthrough.

film_review

Modify a plan not yet committed to the production graph; cannot overwrite an executed production graph.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer)、plan (object).Optional: duration、workspace_id.

Writes state, submits tasks, or saves local files; check authorization first. Exact contract: MCP inputSchema · film_review Returns the corresponding walkthrough.

film_quote

Save the latest quote for the remaining work without submitting generation.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer).Optional: targets、video_resolution、shot_duration、reference_mode、export_resolution、ratio、workspace_id.

Writes state, submits tasks, or saves local files; check authorization first. Exact contract: MCP inputSchema · film_quote Returns the corresponding walkthrough.

film_confirm

After showing an accurate quote and obtaining explicit session consent, confirm with delegated, the original quote fingerprint, and quota; the web entry remains as a compatibility method.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer)、quote_id (string)、authorized_max_points (number).Optional: confirmation、quote_fingerprint、workspace_id.

Will write state, submit tasks, or save local files; verify authorization first. Exact contract: MCP inputSchema · film_confirm Returns the corresponding walkthrough.

film_control

Pause, continue, update, retry, restore, or rebuild the first cut according to valid state; rebuilding discards manual edits.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer)、action (string).Optional: workspace_id.

Will write state, submit tasks, or save local files; verify authorization first. Exact contract: MCP inputSchema · film_control Returns the corresponding walkthrough.

film_pause

Pauses subsequent scheduling; submitted tasks may still complete and settle.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer).Optional: workspace_id.

Will write state, submit tasks, or save local files; verify authorization first. Exact contract: MCP inputSchema · film_pause Returns the corresponding walkthrough.

film_cancel

Ends production without resuming, keeps results, and does not cancel submitted subtasks.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer).Optional: workspace_id.

Will write state, submit tasks, or save local files; verify authorization first. Exact contract: MCP inputSchema · film_cancel Returns the corresponding walkthrough.

film_shot_preview

Pause and wait for in-flight work to finish, then preview single-shot changes and affected steps without saving or generating.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer)、document_revision (integer)、production_revision (integer)、shot_key (string)、patch (object).Optional: workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · film_shot_preview Returns the corresponding walkthrough.

film_shot_apply

Save the previewed single-shot changes and pause, keeping old results; update and recheck pricing before continuing.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer)、document_revision (integer)、production_revision (integer)、shot_key (string)、patch (object)、impact_token (string).Optional: workspace_id.

Writes state, submits a task, or saves a local file; check authorization first. Exact contract: MCP inputSchema · film_shot_apply Returns the corresponding walkthrough.

film_media_add

Adds existing assets you have rights to use to a specified film without calling a model.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer)、asset_ids (array).Optional: workspace_id.

Writes state, submits a task, or saves a local file; check authorization first. Exact contract: MCP inputSchema · film_media_add Returns the corresponding walkthrough.

film_editor_preview

Previews bounded changes to a specified film, such as basic trimming, ordering, subtitles, and volume.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer)、command (object).Optional: workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · film_editor_preview Returns the corresponding walkthrough.

film_editor_apply

Applies basic edits to a specified film as previewed, without overwriting the whole editor.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer)、command (object)、impact_token (string).Optional: workspace_id.

Writes state, submits a task, or saves a local file; check authorization first. Exact contract: MCP inputSchema · film_editor_apply Returns the corresponding walkthrough.

film_export_preview

Checks a specified film's edits and media and freezes the export intent without submitting.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer).Optional: workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · film_export_preview Returns the corresponding walkthrough.

film_export_submit

Submits the previewed local render; replaying the original request does not export again, and past results are kept.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer)、impact_token (string).Optional: workspace_id.

Writes state, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · film_export_submit Returns the corresponding walkthrough.

film_archive_preview

Preview hiding or restoring nodes in a canceled plan; retain assets and history, and handle cross-plan links first.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer)、document_revision (integer)、production_revision (integer)、archived (boolean).Optional: workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · film_archive_preview Returns the corresponding walkthrough.

film_archive_apply

Archive node visibility for a canceled plan as previewed; archived=false restores it, without creating a second plan.

Required: project_id (string)、film_id (string)、request_id (string)、revision (integer)、document_revision (integer)、production_revision (integer)、archived (boolean)、impact_token (string).Optional: workspace_id.

Writes state, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · film_archive_apply Returns the corresponding walkthrough.

project_context_get

Read a bounded overview of the current project, revisions, films, and recent jobs; does not read chats from other clients.

Required: project_id (string).Optional: workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · project_context_get Returns the corresponding walkthrough.

project_images_get

View image thumbnails for up to four nodes you can access; this is not the same as watching the full video.

Required: project_id (string)、node_ids (array).Optional: workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · project_images_get Returns the corresponding walkthrough.

session_capabilities

Read the capabilities the host actually declares; unknown does not mean support for players, attachments, or embedded components.

Required: none. Optional: none.

Read-only / preview operation. Exact contract: MCP inputSchema · session_capabilities Returns the corresponding walkthrough.

media_import

Import restricted attachment file parameters supplied by the host, reusing file validation and upload receipts; arbitrary URLs are not accepted.

Required: project_id (string)、request_id (string)、file (object).Optional: category、canvas_reference、workspace_id.

Writes status, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · media_import Returns the corresponding walkthrough.

media_upload_prepare

Issues a five-minute upload entry point for clients with local file permissions, without letting remote servers read paths on your computer.

Required: project_id (string)、request_id (string)、file_name (string).Optional: category、canvas_reference、workspace_id.

Writes status, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · media_upload_prepare Returns the corresponding walkthrough.

media_delivery

Returns asset metadata, a limited thumbnail, and a five-minute direct download link; revoking access takes effect immediately, with no need to open the editor.

Required: project_id (string)、asset_id (string).Optional: preview、workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · media_delivery Returns the corresponding walkthrough.

media_adoption_preview

Previews replacing the original generation node with an externally uploaded asset or a historical candidate, lists downstream effects, and does not create a new plan.

Required: project_id (string)、request_id (string)、revision (integer)、production_revision (integer)、node_id (string)、asset_id (string).Optional: job_id、workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · media_adoption_preview Returns the corresponding walkthrough.

media_adoption_apply

Adopts the asset using both revisions, the impact token, and the original request; keeps history, reuses the original step, and recalculates pricing after updating downstream.

Required: project_id (string)、request_id (string)、revision (integer)、production_revision (integer)、node_id (string)、asset_id (string)、impact_token (string).Optional: job_id、workspace_id.

Writes status, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · media_adoption_apply Returns the corresponding walkthrough.

object_trash_list

Queries the trash in the current workspace; restoring requires resolving the parent work and episodes first.

Required: none.Optional: workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · object_trash_list Returns the corresponding walkthrough.

object_get

Reads the version, trash status, references, and active jobs for a work, episode/canvas, or asset.

Required: oid (string)、kind (string).Optional: workspace_id.

Read-only / preview operation. Exact contract: MCP inputSchema · object_get Returns the corresponding walkthrough.

object_preview

Preview renaming, asset categorization, soft recycle, or restore; re-check if the version changes.

Required: kind (string)、oid (string)、request_id (string)、version (string)、action (string).Optional: value、workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · object_preview Returns the corresponding walkthrough.

object_apply

Apply the original preview; the receipt covers recycle and restore, and does not permanently delete media.

Required: kind (string)、oid (string)、request_id (string)、version (string)、action (string)、impact_token (string).Optional: value、workspace_id.

Writes state, submits jobs, or saves local files; verify authorization first. Exact contract: MCP inputSchema · object_apply Returns the corresponding walkthrough.

browser_open

Upload assets on the current project webpage, or view and download the film; after completion, read back the stable asset ID.

Required: project_id (string)、purpose (string).Optional: film_id、workspace_id.

Read-only/preview operation. Exact contract: MCP inputSchema · browser_open Returns the corresponding walkthrough.