

A TWO-MINUTE FILM / MAKING OF

The Last *Lantern*

A hyper-real short: a traveller carries the forest's last lantern to a door in a great tree. The cameras, the cuts and every position were set in a Blender grey-box first. Scenario models made the set, the stills, the 40 shots and the sound, each one held to that layout.

2:00

40 cuts, 24 fps

43

generated 3D models

39


Seedance takes

29,342

credits (CU)

One agent, nine Scenario models, one layout.

Claude Opus 5.5 in Claude Code, connected to the Scenario MCP, chose and compared the models, ran the generations and built the set in Blender 5.2 through the Scenario Blender Plugin. Long batches such as the video takes ran on the same models through the Scenario API. The director gave the brief, picked between options and sent notes by cut and frame.



The final film, in credits

39 video takes, Seedance 2.5	24,694 CU
39 stills	1,030 CU
43 models, Tripo 3.1	2,580 CU
43 concepts	129 CU
6 reference sheets	158 CU
3 ground materials	39 CU
1 sky	44 CU
Final sound	668 CU
Total	29,342 CU

What every generation in the final film cost. Blender renders, edits and checks cost 0 CU.

A skill to do it again

The method is packaged as an agent skill, Scenario MCP skill: Blockout to Film, free on the exploration page. Give it your story in any style and your agent runs the same steps with the same checks.

Every generation went through one agent and the same layout; Blender renders and checks cost 0 CU.

Built in grey before any picture.

The brief came in two messages: "Can you create a block out scene for an incredible enchanted forest with lots of various trees, rocks, arches, magical artifacts and elements, some relief (ground is not even) etc. About 50 m by 50 m", then "How can we use this scene in the Scenario MCP? To write a 2 min mini film". Claude Opus 5.5 built the scene in Blender 5.2 through the Scenario Blender Plugin.



The build in stage order: terrain, water, the great tree.



The finished blockout: 289 objects.



Same frame in wireframe.

One colour per kind of object

Uneven terrain, about 120 trees and bushes, 59 rocks, five arches and a bridge, 78 magical props. Each kind has its own flat colour, so a model can be told later what every shape is. Ten staged scripts build it from an empty scene; the how-to film shows the objects arriving in the order of those stages.

Numbers before pictures

A numeric audit closed the stage: 14.0 s of frame math, no floating objects, both bridge ends on the banks, no overlaps between proxies, the walk camera never closer than 0.45 m to a surface. Then a 14-shot film plan, a previs and an animatic, all at 0 CU.

43 models, each in its block's place.

Each kind of block became a concept image, then a 3D model, then took the block's place at the same position and scale. Proxies were hidden, never deleted, so any model can be checked against the block it replaced.



43 concepts, GPT Image 2.5 Sunburst, 3 CU each.



Blockout, one flat colour per kind.



Same camera, generated models.

Image first, then mesh

Tripot 3.1 in Smart Low Poly mode made every model from its concept, at 60 CU each; for the traveller, Meshy 7.1 ran on the same concept and came close, so Tripot was kept. PATINA Material made the forest floor, trail and creek bank, and Scenario Skybox GPT the dusk clearing around the set.

Fitted to the block

Bushes and boulders first sat 40 to 55 % under the ground, because their proxies were half-buried spheres. The placement script now fits each model to the visible height of its block; floating props keep their height and the creek log keeps its span.

From a grey block to a 3D model.

Every model starts as a block in the layout. Each kind of block gets a concept image, the concept becomes a textured mesh, and the mesh takes the block's exact place. Two of the 43, the elder tree and a lantern post:



Elder tree: block, Blender



Concept, GPT Image 2.5 Sunburst, 3 CU



Model, Tripo 3.1, 60 CU

The block decides where

Position, scale, footprint and the flat colour that names its kind. The model is fitted to the block's visible height and takes its place; the block is hidden, never deleted.

The concept decides what

Shape, material and detail. When a model does not suit its role, it is fixed in Blender: the crystal shards around the altar stood on a mossy mound, so they were cut at 24 % of their height and mirrored into floating shards.

Kept light


Tripo 3.1 in Smart Low Poly mode, 2,000 to 12,000 faces per model. A fern that came back at 29,000 faces was decimated to 3,500.



Lantern post: block, Blender



Concept, GPT Image 2.5 Sunburst, 3 CU



Model, Tripo 3.1, 60 CU

A clearing under a painted dome.

The set stands inside a 62 m dome that carries one panorama by Scenario Skybox GPT: a round clearing at dusk, seen from eye level, 3840 x 1920, 44 CU. The same image lights the scene, so the set takes the colour of the sky around it.



The skybox, unrolled.



The dome from outside, near half culled.



Inside, in Material Preview.

The horizon lands on the dome's foot

Measured on the image, the tree bases sit 18.7 degrees below eye level. The panorama is projected from a point 21 m above the floor (62 m x tan 18.7 degrees), so the painted trees meet the ground where the dome begins. The ground continues to the dome, and a ring of about 250 generated trees, 9 to 17 m tall, gives wide shots their parallax.

One prompt, free to download


The prompt asked for the same wall of trees at the same distance in every direction, and "no people, no buildings, no paths, no lights, no text". A first sky seen from above left the set floating in mid-air; the eye-level clearing fixed it. The skybox is free to download for your own scenes.

[Download the skybox, JPEG 3840 x 1920](#)

[Scenario Skybox GPT](#)

Every light has a source.


A first rig with a sun, canopy shafts and drifting fireflies was turned down in review: light comes only from things in the story, each in its own colour, nothing random.



Cut 10, clay: the arch's canopy keeps its teal.



Cut 10, Eevee look.



Cut 31, clay: the portal ring, grey.



Cut 31, look: its runes wake after the orb.

Posts, lamps, crystals, wisps

The second rig: 7 lantern posts, 16 trail lamps, 71 crystals in six colours, spirit-tree leaves, glow mushrooms, 28 wisps, rune stones, and a portal with six orbiting orbs. Every glowing surface shares its light's keys and noise, so a candle flicker shows on the flame and on the light it casts.

Light on cue

The forest sleeps until the orb is relit. Then a wave runs out from the altar at 2.88 m/s, each lamp wakes within 1.5 to 4 m of her lantern and fades behind her, and from frame 2790 every light turns gold for the finale. These rules apply at render time; the scene files stay untouched.

40 cuts, planned to the frame.

Every beat keeps its main camera and gets coverage cameras on the same action at the same frames, so a cut inside a beat matches by construction. The edit list sets the first and last frame of each cut: 40 cuts from 2.0 to 5.5 s, 3.0 s on average.



Top view: the camera of each of the 39 takes, coloured by act, and her path, dotted.

Framing checked by rays

A ray-cast audit checks a coverage camera in seconds: the lens against the geometry, the line to the aim point, her feet, chest, face or flame in frame and seen. Round one found 11 cameras inside trees or rocks or missing her; after three rounds none failed. Fixes moved the blockers or the camera, never the plan.

Cameras as specs

Coverage cameras are written as specs, relative to her, to the ground or absolute, and built in memory at render time. Nothing is saved into the scene, so one change to the plan re-renders every cut that uses it.

The same frame, four times.

For each cut, one frame goes through four layers: the clay playblast (grey set, light sources in colour, a capsule for her), the EEVEE look, the finished still and the frame of the film. They line up because each one is made from the one before. Here, cut 9 at frame 648.



Clay playblast, Blender Workbench



EEVEE look



Still, GPT Image 2.5 Sunburst



Film, Seedance 2.5

From render to photograph

GPT Image 2.5 Sunburst at xhigh, 2560 px (26 CU), turns the look frame into a still. The prompt keeps the exact composition, lens, horizon and every silhouette, and calls the render's textures placeholders, to be rebuilt as photographed detail. On a test frame, xhigh was clearly sharper than high; max showed no visible gain.

One still per cut

With stills from several moments of a shot, Seedance cut or dissolved to reach each one, even when told not to. With the clay video and one still from inside the cut, every take stayed one continuous shot.

One lantern, one door, one face.

In the first cut, the lantern changed from shot to shot and each door shot invented its own runes. The fix was a reference sheet for each recurring element, made from the 3D model itself, so the sheet matches the set.



Traveller, from her two largest finished views.



Lantern, from a turnaround of the mesh.



Door: front, three-quarter, knob, runes lit.

From the mesh, then to picture

Turnarounds of the lantern and the door were rendered from their meshes (0 CU), then GPT Image 2.5 Sunburst made each sheet for 26 or 27 CU. A pose sheet shows the lantern lowered, raised, and offered in both hands under the base, so the flame stays clear of her fingers.

Sheets go to the stills only

Each still that shows the lantern or the door carries its sheet, with a clause that keeps the size, position and crop of the layout; a test with the in-hand sheet and no such clause zoomed out. Seedance never sees a sheet: it gets the clay and one still, and the still already carries the right lantern.

Seedance films the clay.

Seedance 2.5 runs in reference mode: the cut's clay video as @video1, one finished still as @image1, and a prompt that says which input governs what. Each take covers its cut plus a few frames of handle, at 1080p, and the edit trims it to its frames.



Cut 21, clay: the capsule stand-in on the hill path.

A capsule, not a statue

A rigid statue of her slid along the path, and so did her takes; a rigged walk gave sliding feet. A legless capsule with a ball head worked: the prompt says it only marks where she is, her height and her direction, and Seedance supplies a natural walk.



Cut 21, Seedance 2.5 take.

What the prompt pins down

"The framing always comes from @video1." Every prompt also asks for one continuous take with no cuts or dissolves, states the story state (the door stays dark, the lantern warm amber), and names a prop's colour when the still does not show it. Each take is checked against its clay: overlay, cut and dissolve scans, a dense contact sheet, her feet.

Notes by cut and frame.

The edit list assembles the 40 cuts from their takes. Each review watched the film over its playblast, with cut and frame numbers, so every note named a cut and a frame, and every fix touched that cut alone.



Cut 35 before: a second lantern by the mushroom.



After: one lantern, in her right hand.



Cut 21 before: her pace was too slow.



After: the clay plays 1.25 x faster.

Four cuts in one day

The first cut came twelve hours after the first message. Three review rounds followed: reference sheets and a reordered door sequence, then single-cut fixes such as a longer stride on the hill and one lantern instead of two in cut 35.

Checks that stop regressions

At every cut boundary a pair sheet checks that no cut steps back in time and that her lantern is warm up to frame 2065 and dark after.

The sound follows the cuts.

The narration is cut into 24 phrases on its word times, each placed on the cut that shows it: the lantern on its close-up, the arch on the arch, the door on its opening. The score follows the story in 14 sections, and its accents meet the orb, the wave and the door.



The door opens onto the golden forest.

Voice, score, effects

The narrator is Bill on ElevenLabs v4, chosen by ear from 19 setups on four voice models, each first ranked on pitch movement, pace and fit on the cuts. The score is ElevenLabs Music Advanced v2.5, written section by section from a 14-section brief at 120 BPM, from the night to the open door, and checked with whisper for stray voices. Twelve ElevenLabs SFX 2 hits sit on their events, over Sonilo V1.1 foley generated from the picture.



The last cut: every light turned gold.

The receipt

Seedance 2.5, 39 takes	24,694 CU
Stills, 39	1,030 CU
3D set: 43 concepts, 43 models, 3 materials, 1 sky	2,792 CU
Sound: narration, score, effects, foley	668 CU
Reference sheets, 6	158 CU
Total	29,342 CU

Every generation in the final film, from the project's job ledgers. Blender renders, edits and checks cost 0 CU.