

GodwinMix: scope, timeline and cost

How an engagement works, what you supply, what hardware to plan for, the licence, and the risks you accept by choosing a young product.

OCTOBER 2026

What you are paying for

GodwinMix itself is free. It is open source under the Apache License, Version 2.0 (January 2004), as stated in the project's LICENSE file. There is no hosted service and no account: it is a program you run on your own machine. There is no licence fee to us or to anyone else for the mixer.

What you pay us for is the work around it: choosing and applying the right preset, setting up your sources and destinations, testing on your actual machine and network, training the people who will run it, and, where you need it, building something new such as a plugin, a graphic package or a custom branded build.

Some parts that travel with it carry their own terms. The Windows installer includes the NDI runtime on the NDI SDK's terms, and GStreamer and Chromium licence files travel with the installers. If you use an AI agent to direct a show, the model costs are paid by you to the AI vendor.

How an engagement works

The size of the job depends far more on your show than on the software. Three typical sizes, all estimates:

Engagement	Typical example	Estimated time	What happens
An afternoon	A church already streaming from a laptop	Half a day on site or remotely	Apply the church preset, put in your stream keys, connect your cameras, add phones as extra cameras, set up lyrics over the picture, stream a test to an unlisted event, show your volunteer the red button.
About a week	A classroom, a sports club, a conference	3 to 5 working days, spread over two weeks	Apply the nearest preset, fit it to your rooms and kit, set up recording, graphics and a score or data feed if needed, run two or three rehearsals, train two or three operators, write a one page run sheet for your people.
A month or more	A broadcaster's contribution chain, a 24/7 AI run channel, a bespoke plugin	4 weeks and up	Design the chain with your engineers, test SRT links and tally with your equipment, run long soak tests on your hardware, build and test any plugin you need, agree what happens when something fails at 3 a.m.

For the larger jobs, part of the time is simply waiting and watching: a channel meant to run unattended has to run unattended on your machine for a good while before anyone should trust it.

What you supply

- A machine to run it on. A desktop or laptop for a staffed show, or a Linux server (or a container host) for a headless one.
- Your cameras, or phones to use as cameras.
- Your destinations and their stream keys (YouTube, Facebook, a broadcaster's SRT address and so on). We never need to see a key after it is entered; the interface does not show it again.
- A network that can carry it. Upload speed is the usual limit. Four 1080p60 NDI feeds need about 500 Mbit on a network of their own.
- For remote guests on other networks, a way to serve the mixer over https: a domain name and a reverse proxy, which we can set up with your IT team.
- People to train, and time for rehearsals.

Hardware guidance

A GPU is optional. The same program picks hardware codecs at startup if they are there and falls back to software if they are not.

Platform	Hardware decode and encode	Notes
Windows	Direct3D 11 decode, Media Foundation encode; NVIDIA if present	Quick Sync was picked by itself on the Intel test laptop. On that laptop the driver corrupted memory on about one show start in four; the mixer restarts the show.
macOS	VideoToolbox	Installer is for Apple Silicon.
Linux	NVIDIA, or VA for Intel and AMD	Best fit for a headless server.
Any	Software (x264)	Works everywhere, uses more processor.

On an Apple M4 Pro, two 720p30 sources with encoding used 0.70 of one core with hardware encoding and 0.88 without. Smaller machines (Raspberry Pi, Intel N100, a plain laptop) have not been measured. The honest guidance is: try it on the machine you plan to use, which is the first thing we do in any engagement.

Some practical rules from the presets: 1080p60 with no GPU is more than a laptop can do, so drop to 30 frames a second or use a machine with a GPU. A 720p canvas is kind to a room laptop on battery.

Our fees

Item	Fee
An afternoon setup (preset, sources, destinations, a test stream)	[fee to be confirmed]

Item	Fee
A one week engagement with rehearsals and training	[fee to be confirmed]
A month or more: contribution chain, unattended channel, soak testing	[quoted per project]
A bespoke plugin	[quoted per plugin]
A custom branded build under your product name	[fee to be confirmed]
Support during your live shows, or a monthly retainer	[monthly fee to be confirmed]

Ways to pay

For the smaller engagements, one fee on completion. For a month or more, a fee split across agreed milestones, or spread monthly together with ongoing support. [Payment terms to be confirmed.]

Honest scope

What GodwinMix does not do today:

Not available	What it means for you
Proven production parity with OBS, vMix or Wirecast	The project does not claim it. Endurance acceptance (one hour and cross platform) is still outstanding.
Unattended use, recommended	Not until that endurance testing finishes.
Instant replay	No replay plugin exists, though the esports preset names one.
Windows game capture	Not in the first year. Screen capture is a whole display or a region.
Vertical and horizontal output together	One picture shape at a time.
A hosted service	You run it. There is no account.
Per guest invite links	A guest link carries the channel's key until invites exist.
Stinger sound	A stinger's own sound is dropped today.
Moving graphics, transparency and transition packs on macOS and Linux	Run on Windows only so far.
Some media pieces in the Linux installer	It lacks <code>cmafmux</code> , <code>livesync</code> , WHIP and WHEP elements and AV1 decoding; HLS output from the Linux installer is affected. Docker is the safer route on Linux for now.
WebP pictures on Windows	Do not open.
DeckLink SDI and HDMI capture	In the code, never run against a real card.
NDI from a real third party sender	Not yet tested.
Published performance on small machines	Not measured.

Not available	What it means for you
A plugin marketplace	Not yet.

Several October 2026 features (Studio mode's new take bar, the suggested preview, transition packs, the Graphics gallery, moving HTML graphics and one button AI setup) are on the main line of the code but not in a tagged release. Some fixes, including a large cut in preview delay for a late camera and more complete installers, are described on working branches dated 6 October 2026; we will confirm which release carries them before your engagement starts.

Risks, and who decides

Risk	What we do	The decision you make
A long unattended run fails in a way not yet seen	Run a soak test on your machine before go live; read the session log; set up alerts	Whether to run unattended now, or keep someone watching until endurance acceptance is done
The mixer stops responding after many hours	Calls now time out after five seconds and report what is holding the loop; we monitor <code>/api/status</code>	Whether to schedule a restart between shows
A feature you need has only run on Windows	Test it on your platform first, in the first days	Whether to run on Windows, wait, or do without
The installer for your platform lacks a piece	Use Docker or a full GStreamer install on Linux; check the release notes	Which deployment route you accept
A guest link is passed on	Tell you it carries the channel key; change the key after the event	Whether remote guests are acceptable before invites exist
Your uplink is too weak	Measure it first; lower the bitrate; the outage buffer holds 5 seconds	Whether to upgrade the connection
Hardware encoder driver faults	Choose the codec deliberately; fall back to software	Which machine and codec to standardise on
AI agent costs grow	Use small snapshots (84 tokens a frame) and the smallest tool profile	Which model, how often it looks, and a monthly budget
The project changes direction	<code>api_level 1</code> is frozen for breaking changes, with a year of overlap after any break; the code is Apache 2.0 and yours to keep	How much you build on the API now