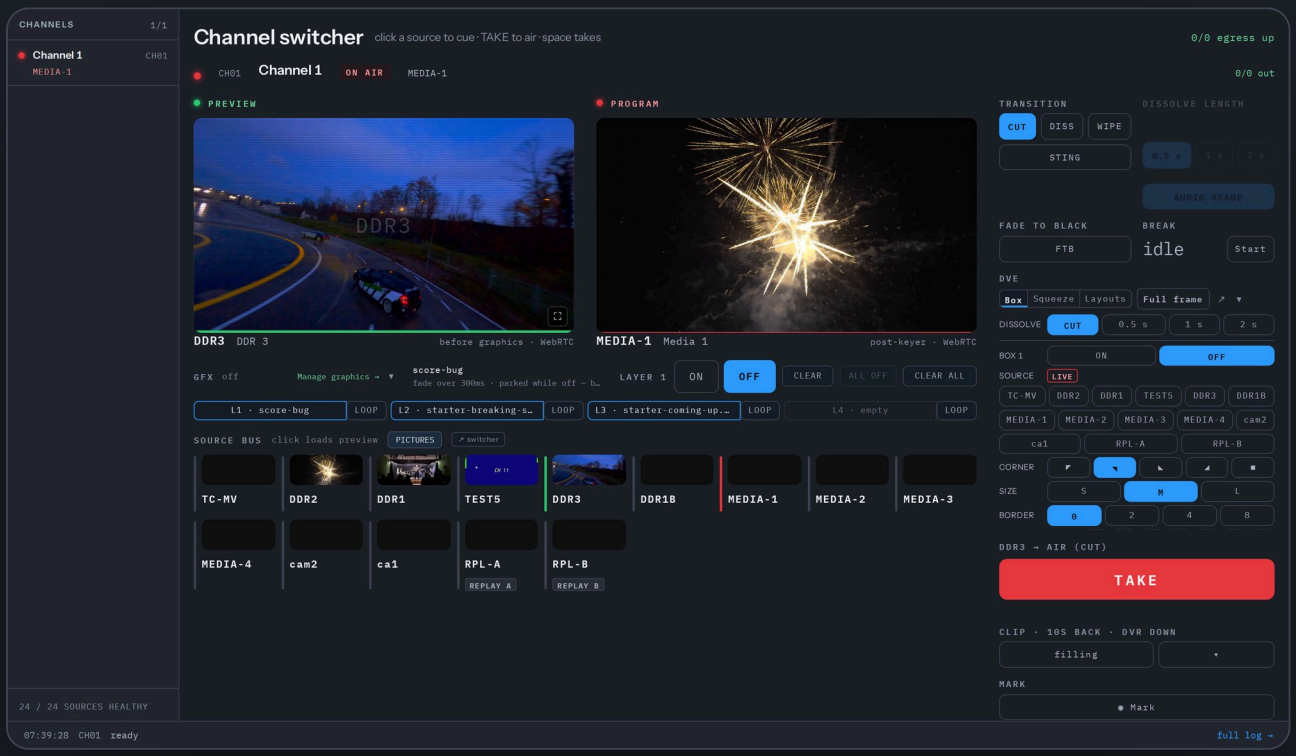
Feature guide

One system for live production, master control and FAST channels — with captions, translation, AI voices, comms and replay built in, run from a web browser.

ON AIR PREVIEW CAUTION ARMED



Contents

This edition describes MCR as it stands in November 2026. Anything still being proven is marked BETA, anything that needs extra set-up is marked OPTIONAL, and anything that runs only on an NVIDIA graphics card is marked NVIDIA GPU.

00	**A complete production system**	What MCR is, in one page	3
01	**AI that runs where you choose**	Local or cloud, per service, per channel	4
02	**Where MCR fits**	The jobs one system covers	6
03	**The platform**	Channels, templates, settings and licensing	7
04	**Multiviewer & monitoring**	Walls, monitors, WebRTC and NDI	8
05	**Sources & contribution**	SRT, RIST, RTMP, NDI, WebRTC, guests and PTZ	10
06	**Switching, transitions & macros**	Preview/program, DVE, chroma key, second M/E	12
07	**Graphics & DVE scenes**	Make, Play, the Board, templates, NDI, web pages, live data	13
08	**Audio desk, 5.1 & AC-3**	Strips, buses, dubs, loudness, surround	15
09	**Programme mixes, commentary & comms**	Follow mixes, international clean, party lines	16
10	**Outputs, HLS & delivery presets**	SRT, E-RTMP, HLS ABR, NDI, encoders	17
11	**Captions & translation**	608/708, WebVTT, burn-in, language detection	19
12	**AI voices & audio description**	Dubbing per language, on the desk and on every output	22
13	**Transcripts, the transcript editor & search**	Recordings, library files, feeds and shots	24
14	**Playout, rundown & scheduling**	Decks, rundown, breaks, SCTE-35	26
15	**Traffic, programme guide & FAST delivery**	Day log, BXF, EPG, EIT, platform feeds	28
16	**Recording, clips & highlights**	DVR, ISO, quick clips, highlights	29
17	**Instant replay**	Mark, cue, angles, slow motion, reels	31
18	**Editing, render nodes & publishing**	Editor, house conform, social 9:16, destinations	32
19	**Alarms, hardware & operations**	Alarm policy, hardware verdict, analytics	34
20	**Security & access control**	Logins, roles, MFA, SSO, join codes	36
21	**Automation, API & AI agents**	API v1, Companion, MCP	37
22	**Scale out: nodes, storage & resilience**	Render, transcode, AI, ingest, NDI matrix, HA	38
23	**Formats & standards**	What goes in and what comes out	40
24	**How it behaves, and what is next**	Design rules, upgrades and the roadmap	41
—	**Credits & notices**	Pictures, trademarks, contact	42

Screens are from a demo channel; pictures, names and text in them are examples (credits on the last page).

00 A complete production system

MCR is a full live production and playout system in software. One product does the work of a production switcher, a master control room and a FAST channel playout platform, with live captions, translation, AI voices, crew comms and replay built in. Every channel is run from one web panel, so operators need nothing but a browser, and channels run in the cloud or on your own servers — on ordinary CPUs, with a graphics card added only for the AI you want to run yourself.

Production switcher

LIVE EVENTS AND SPORT

- Preview / program with cut, dissolve, wipes and stingers
- 8-box DVE, chroma key, scenes, macros and a second M/E bank
- Six graphics layers and 23 animated templates
- Instant replay with frame-accurate angles
- Browser guests, PTZ cameras, party lines and talkback

Master control

RUNNING CHANNELS 24/7

- Multiviewer with broadcast layouts, tally and meters
- Primary / backup failover and self-healing inputs
- Loudness to EBU R128 or ATSC A/85, corrected by default
- Alarms with history, as-run log and compliance recording
- Many channels on one box, created from the panel

FAST & linear channels

DELIVERY TO PLATFORMS

- Playout decks and a rundown built from the library
- Traffic, ad breaks, SCTE-35 and HLS cue tags
- House-format checks and conform before air
- Programme guide, EIT and platform EPG exports
- HLS origin with ABR ladders, subtitles and dubs

32
sources per channel — live feeds, media players and replay

6
keyed graphics layers, run from a one-press board

8
DVE boxes per scene, each with its own chroma key

8
replay channels (A–H) on a channel sized for them

28
languages in the live caption catalogue; which are local or cloud is set out on the captions page

608 /708
live captions embedded in the video, plus WebVTT, HLS subtitles and burn-in

5.1
programme bus with AC-3, E-AC-3 and SMPTE 302M delivery

0
software to install for operators or remote guests

HOW A CHANNEL FITS TOGETHER

Sources
SRT · RIST · RTMP · NDI · HLS · WebRTC · browser guests · media players

→

Switcher
preview / program · transitions · DVE boxes · chroma key · macros · PTZ

→

Graphics
6 keyed layers · animated templates · HTML · NDI with alpha · live data

→

Audio & access
desk · buses · loudness · 5.1 · captions · translation · AI voices · AD

→

Outputs
SRT · RTMP(S) · HLS · NDI · recording · clean feed · multiview

AROUND THE AIR PATH recording and replay · clips and the editor · transcripts and search · rundown, traffic and the programme guide · comms · alarms, hardware and the as-run log · API, macros and Stream Deck · render, transcode, ingest, AI and NDI matrix nodes

01

AI that runs where you choose

Speech recognition, translation, voices and the rest each have a local engine and a cloud engine. Run them on your own hardware, or through the providers you already use with your own API keys — chosen per service, per channel.

LOCAL

On your own hardware

Captions	NVIDIA Parakeet or Nemotron streaming, word by word NVIDIA GPU
Speakers	Speaker labels for live captions and for files NVIDIA GPU; sound events such as applause and cheering on the CPU
Translation	Opus-MT, on the CPU or the GPU, including source-to-target pairs such as Spanish → English
Voices	Kokoro text-to-speech, one voice per language NVIDIA GPU
Files	Canary file transcription; summaries, topics and chapters from a local language model NVIDIA GPU
Pictures	Shot type, people and ball for highlights NVIDIA GPU; face tracking for vertical crops on the CPU

With every service local, programme audio never leaves your network, and there is no per-minute bill.

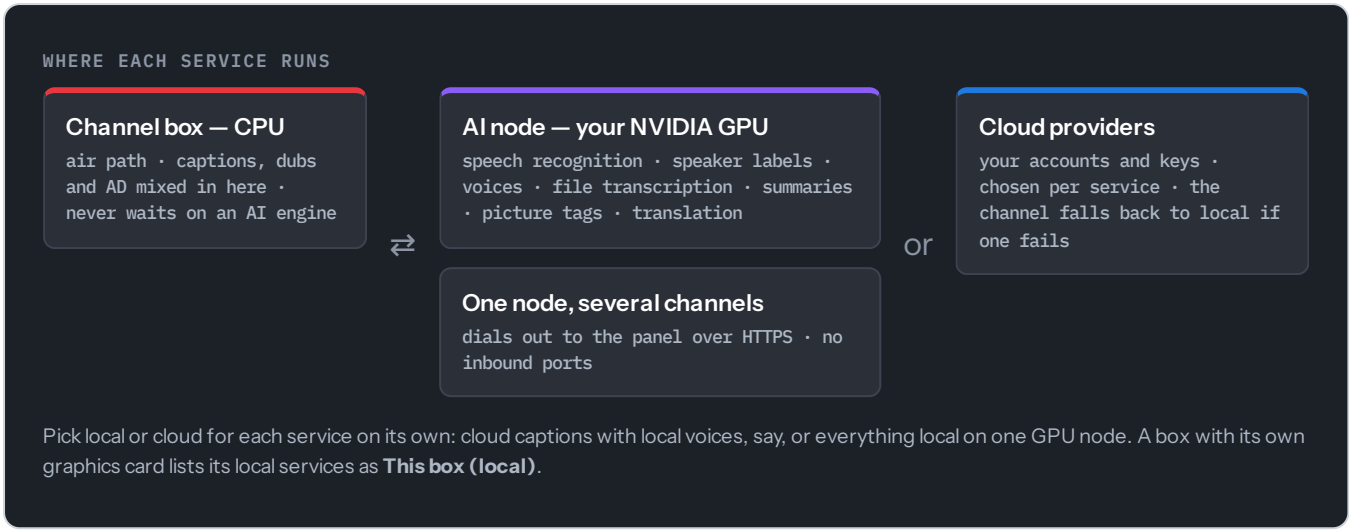
CLOUD

Through your API keys

Captions	Speechmatics (with speaker identification), Mistral Voxtral, OpenAI or Groq Whisper; Deepgram and AssemblyAI BETA
Translation	Mistral, OpenAI-compatible models or Groq
Voices	ElevenLabs, OpenAI or Mistral Voxtral
Description	Mistral vision for AI audio description

Your own accounts and keys, stored write-only in the panel and never shown again. Cloud services need no graphics card at all.

A CPU is enough to run a channel; a GPU is for local AI. Switching, graphics, the audio desk, recording, replay and every output run on ordinary CPUs, and video is encoded on the CPU unless you choose a hardware encoder. Local speech recognition, AI voices, speaker labels, file transcription, summaries and picture tagging run only on an NVIDIA graphics card — on an AI node, or in the channel box itself. There is no slow CPU fall-back for them: a box without the card offers the cloud services instead.



Pick local or cloud for each service on its own: cloud captions with local voices, say, or everything local on one GPU node. A box with its own graphics card lists its local services as **This box (local)**.

A fall-back, not a surprise bill

If a cloud service fails, the channel falls back to the local engine and raises an alarm. Moving from local to cloud happens only if you have switched that on, so an unattended channel does not start sending programme audio off-site or running up a bill by itself.

Air comes first

A slow or failing engine only pauses captions or a dub. The picture and sound on air carry on, and an alarm says what happened.

Voices load when used NVIDIA GPU

A local voice takes no graphics memory until its language is used. It stays loaded through pauses while the language is on, and gives the memory back after 15 minutes idle.

Nothing to switch on by accident

Every dub language starts parked. Nothing is synthesised, and nothing is billed, until someone switches that language on.

The card is shared with care

The Nodes page shows each card service by service: what it holds now, its peak, and the room left, in words such as “room for 1 more caption channel or 3 dub voices”. A dub voice, or a caption channel on the box's own relay, that will not fit is refused before it starts, with what to do instead. File transcription, speaker labels and summaries wait for room rather than push a live caption or voice off the card.

Local captions on several channels NVIDIA GPU

One box or AI node can run up to eight local caption relays, one per channel, as many as its card has room for. A new channel takes the next relay, and one past the limit is refused at creation with the reason.

02

Where MCR fits

Most facilities buy these as separate products, each with its own logins, alarms and support contract. MCR does them in one system, with one panel, one alarm bus, one as-run log and one API.

JOB	USUALLY A SEPARATE PRODUCT, FOR EXAMPLE	IN MCR
Production switching	vMix, TriCaster, Blackmagic ATEM	Switcher with transitions, stingers, 8-box DVE, chroma key, scenes, macros, a second M/E bank and PTZ control
Multiviewer	Hardware multiviewers (Blackmagic, TAG, Evertz)	Browser and NDI walls with broadcast layouts, tally and meters
Master control & playout	Imagine Versio, Grass Valley AMPP, Cinegy Air, PlayBox Neo	Playout decks, rundown, breaks, failover, loudness, alarms and as-run
FAST channel playout & traffic	Amagi, Veset	Traffic log, house-format conform, SCTE-35, programme guide exports and HLS delivery
Graphics	CasparCG, Singular.live, Vizrt	Six keyed layers, animated templates, your own brand and fonts, HTML and NDI graphics, live data
Instant replay	EVS, Vizrt 3Play, vMix Replay	Replay channels, frame-accurate angles, variable speed, high-frame-rate slow motion, reels
Intercom	Clear-Com, Riedel, Unity Intercom	Browser party lines with talk, listen and call, and talkback to guests
Cloud production & clips	Grabyo, LiveU Studio, StreamYard	Browser guests, clips, a timeline editor, highlights and publishing
Captions & access	AI-Media LEXI, ENCO enCaption, Verbit	Live captions, translation, AI voices, audio description and transcripts
Recording & compliance	Separate recorders and loudness loggers	Programme and ISO recording with a loudness compliance record

WHAT ONE SYSTEM SAVES

One panel, one login

Operators learn one tool. Personal logins, roles, two-factor authentication and single sign-on cover switching, playout and delivery alike.

One API

A Stream Deck or an automation system drives switching, graphics, replay, breaks and clips through the same documented API.

One place for faults

A fault anywhere — a lost feed, a failing output, a caption provider, a hot CPU — lands on the same alarm bus, and the as-run says who did what.

Fewer boxes

Switcher, playout, multiviewer, recorder, caption encoder and intercom run as software on the same servers, in the cloud or on your own hardware.

Product names are trademarks of their owners and are given only as examples of each category. MCR is not affiliated with or endorsed by any of them. Every product listed has features of its own that MCR does not; ask us for a side-by-side comparison for your workflow.

03

The platform

One web panel for every channel on a box. Channels are created from the panel from a template, and each one is its own isolated process.

Browser-based control

The switcher, graphics, audio, captions, comms and outputs all run in a standard browser, with no plug-ins. The switcher, the DVE, the audio mixer and the decks each pop out into their own window for a second screen, and any monitor goes full screen.

Isolated channels

Each channel runs in its own process with its own ports, so a fault in one never reaches another. Playout decks come back cued at their position after a restart, and programme is restored.

A restart that says what it will do

Before a restart the panel shows the off-air time the last one took, which outputs and recordings drop, and any hazard — a deck about to reach the end of its clip, a busy CPU, an operator still at work. A red hazard holds the restart until someone ticks **Restart anyway**.

Keyboard-first where it counts

Space takes, Esc cancels, 1–9 fire graphics on the Board, and J / K / L shuttle through recordings. Destructive or full-frame actions take two presses.

Channels from templates

Add a channel from **Channels**: pick a template — Live sport / event, Studio / news, 24/7 playout, Contribution / ingest, Simple single feed, Small box (CPU only) or Custom — and a frame rate, adjust how many feeds, decks and replay channels it gets, and tick the options it needs.

Settings in the panel

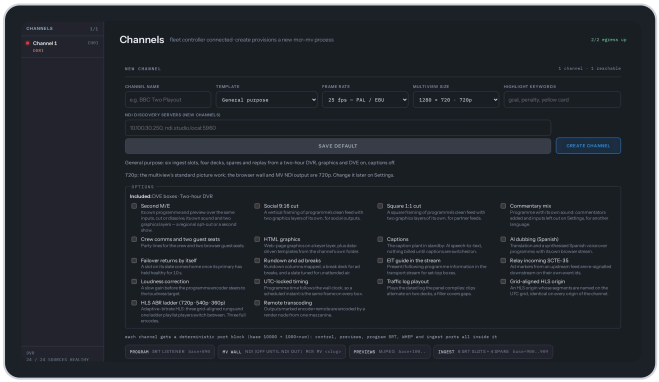
35 settings change on a running channel, and dozens more are staged from the panel and applied at a restart you choose. Export a channel's settings, import them, save named presets, or copy from another channel.

Frame rates

23.976 to 60 fps, including 29.97 and 59.94, chosen per channel, at 720p or 1080p. Replay can also run at 100 or 120 fps for true slow motion.

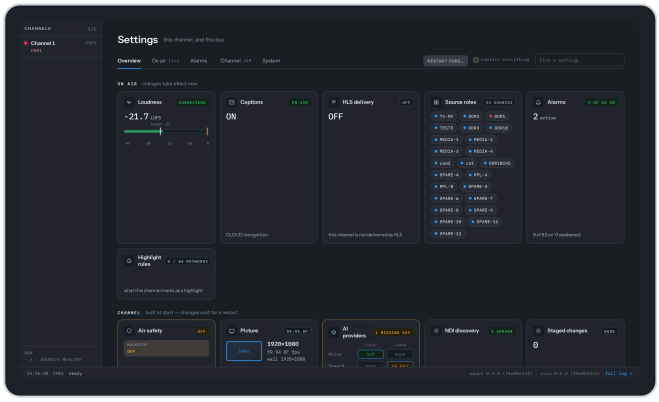
Honest state

If the panel loses contact with a channel, the page greys out and says **NOT LIVE**. You never mistake cached state for the real one.



Channels → New channel. A name, a template and frame rate, the multiview size, then the options a template turns on, each with a line on what it does.

demo channel



Settings → Overview. What is live now — loudness, captions, HLS, source roles, alarms — and what is built at start, with staged changes waiting for a restart.

demo channel

A 30-day trial, then a licence. A new install runs as a 30-day trial: one channel with up to four live sources. The panel counts down the days and warns on every page in the last week. If the trial ends, or a paid licence lapses and its seven-day grace period runs out, the channel's outputs show a **TRIAL ENDED** or **LICENCE EXPIRED** notice instead of programme until a key is entered; entering the key restores air at once, with no restart. The licence agreement sets out the terms.

04

Multiviewer & monitoring

See every source at once, with tally, labels and meters — in the browser, on a gallery screen over NDI, or as an output of its own.

Broadcast layouts

1, 4, 9 or 16 tiles or automatic, plus layouts with a big programme and preview: 2+8, 2+16, 1+5, 1+7 and programme only. Red tally on programme, green on preview, a label and audio meters on every tile. Pictures keep their shape: a 16:9 source in a box of another shape is fitted with black around it, never stretched.

Every source shown

A wall grows to the next layout rather than silently leave an input off, unless it is set to a fixed shape for a physical monitor wall. Saved tile choices follow each input by name.

Browser wall over WebRTC

The multiview and the switcher's preview and programme monitors stream to the browser over WebRTC (WHEP), at under a second, falling back to a slower stream where WebRTC is blocked. **BROWSER** switches the browser wall on or off live; off costs nothing.

NDI out to the gallery

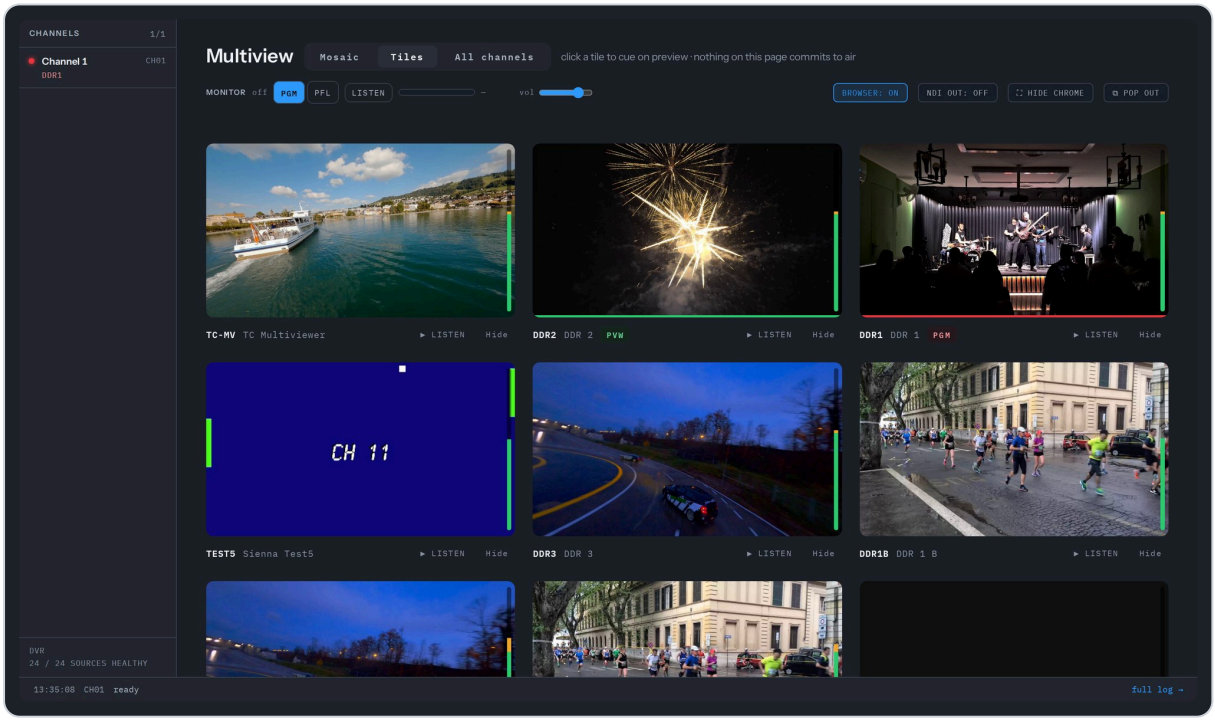
NDI OUT sends the wall to a screen on your network. The browser wall and the NDI wall can be laid out differently OPTIONAL, and the wall can also be an output of its own over SRT, RTMP or NDI, or recorded.

Monitoring quality

Low, Standard or High presets for the browser monitors, each stating what it costs, and a choice of CPU, NVIDIA or Intel for their encoding. Monitors follow the channel's own size and frame rate.

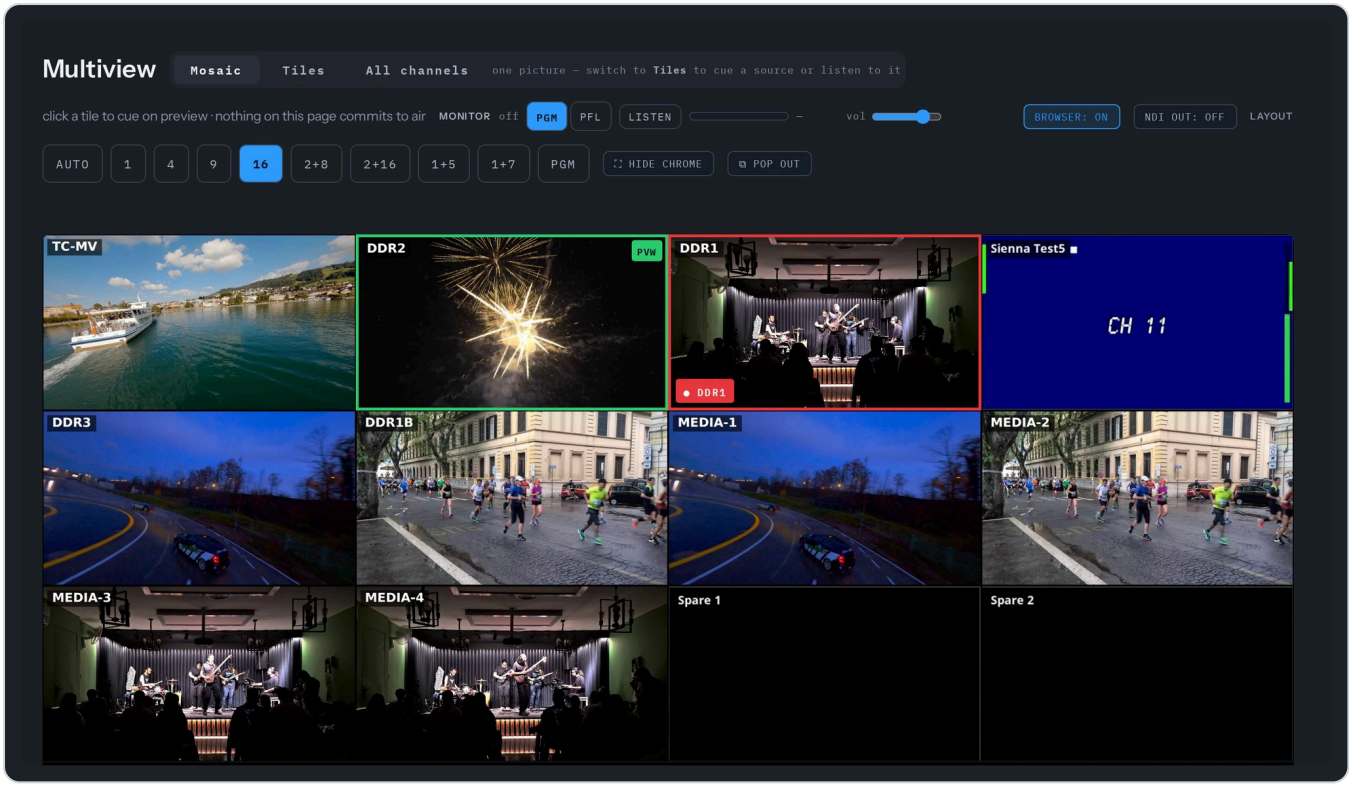
Listen, PFL and every channel

Listen to programme or pre-fade-listen to any source in the browser without touching the mix. **All channels** shows every channel's programme on one page. A monitor that stops says **NO PICTURE** rather than freezing.



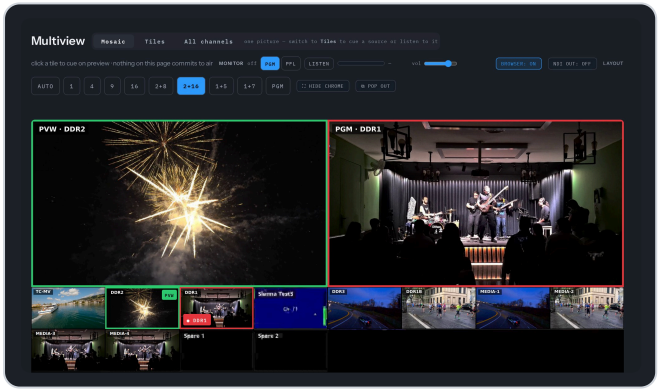
Multiview · Tiles. Every source as its own picture over WebRTC, with its label, tally and audio meter; cue or listen from any tile, with BROWSER and NDI OUT at the top. Empty inputs stay black.

example pictures · credits at the back



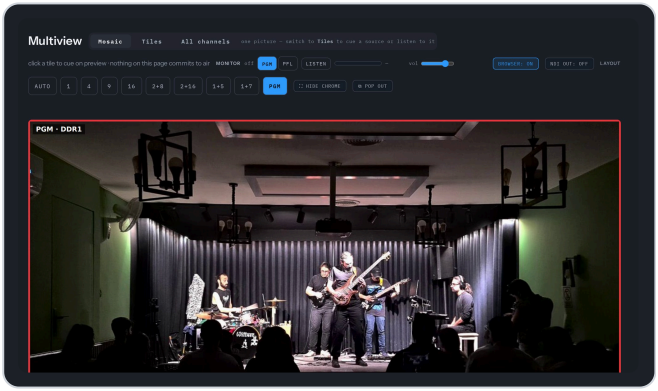
Multiview · Mosaic, 16. One picture composed by the channel itself — the view for a wall screen, a pop-out or NDI — with preview edged green, programme red and its name on a pill; the layout picker above runs from 1 to 16, the broadcast layouts and programme only.

example pictures · credits at the back



Mosaic, 2+16. Preview and programme big across the top, every source below.

example pictures



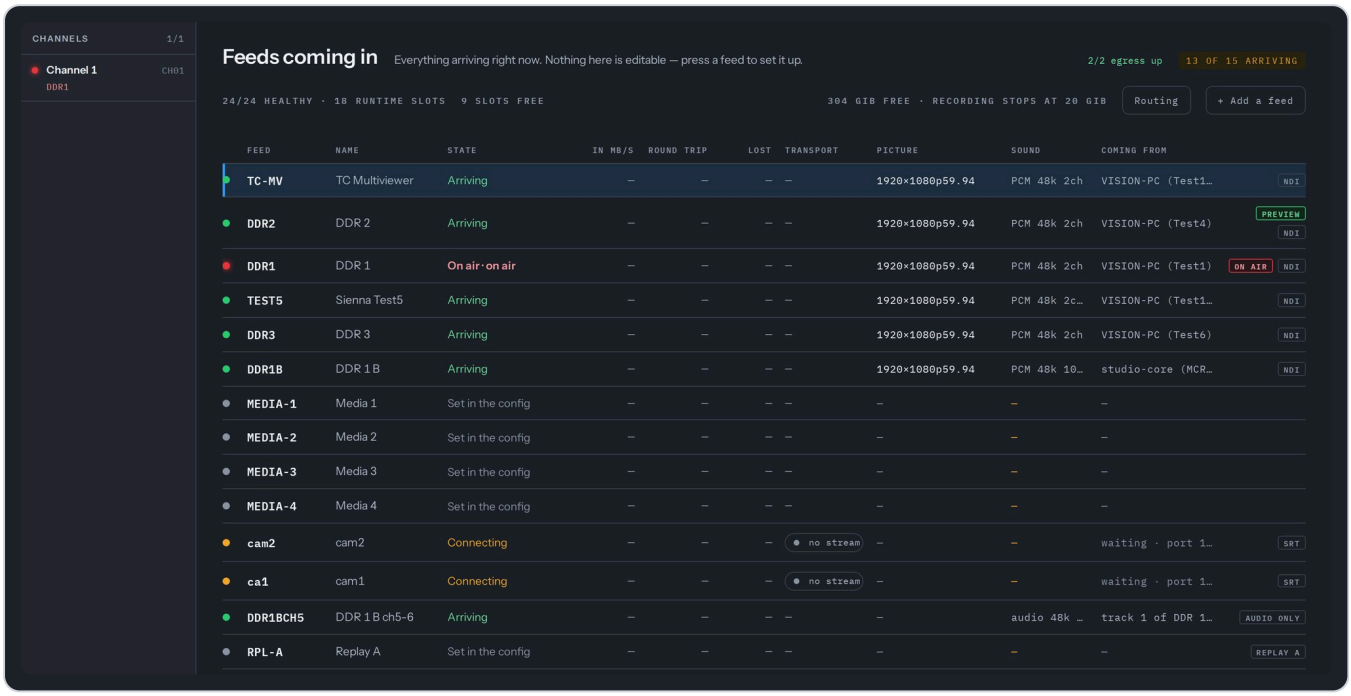
Mosaic, programme only. The whole wall given to what is on air, for a confidence screen.

example pictures

05

Sources & contribution

Bring feeds in over the protocols contribution actually uses. A source is set up once on the Inputs page and appears on the switcher with a live picture. Rename it at any time: the new name shows everywhere at once and the feed keeps running.



Inputs. Every feed with its state, picture and sound format and where it comes from — NDI sources, SRT listeners, media players, an audio track taken from another feed, and replay.

demo channel · example names

SRT and RIST

Listener or caller, encrypted. H.264, HEVC and MPEG-2 video; AV1 where the decoder is installed. Bitrate, round-trip and loss per feed, with TR 101 290 priority 1 and 2 checks. A caller that restarts is taken back within seconds.

NDI, HLS and WebRTC

NDI sources found on the network or through an NDI Discovery Server, with NDI Access Manager groups and transports; HLS pulls; and WHEP/WebRTC ingest that reconnects by itself.

Remote guests from a browser

Send a link and up to eight guests join from a laptop or phone. They wait until you admit them, hear programme without their own voice, and see what you choose. A guest can press **Talk to producer** or **Cough**.

Primary / backup failover

Give a source an ordered backup list — a second feed, then a slate. If the primary fails the channel moves down the list. Returning is a policy you set: automatically, only from the slate, or by hand.

RTMP and Enhanced RTMP

Take RTMP, RTMPS and Enhanced RTMP (HEVC and AV1) from encoder apps such as OBS, vMix and Larix, with several audio tracks and SCTE-35 cues carried through.

Pick the sound you need

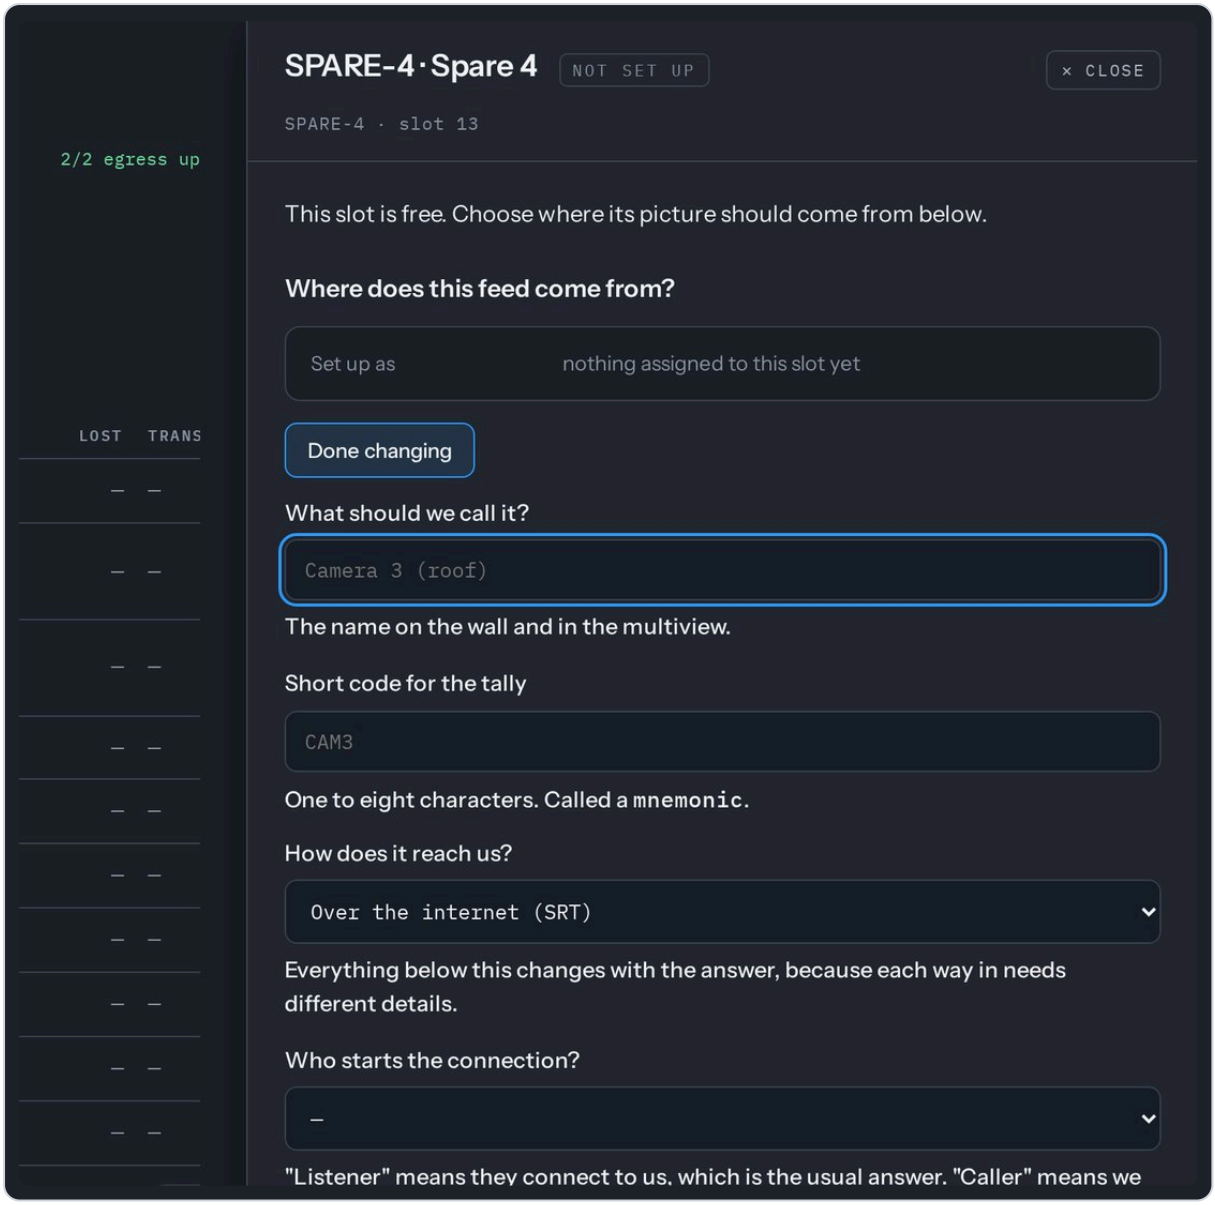
Choose which channels or track of a feed a source takes — a pair, mono, split or 5.1 — and change it live without dropping picture. Each track of a multichannel feed, such as a 16-channel NDI feed, can be its own strip on the desk.

PTZ cameras BETA

Pan, tilt, zoom, speed, home and six presets for the camera on preview, over VISCA over IP or TCP. Cameras are set up on Inputs, and the PTZ controls appear under the preview only when the camera on preview has them.

Decode where it suits you

Decode each feed on the CPU, an NVIDIA card or Intel Quick Sync, or let the channel choose. A feed a graphics card cannot decode falls back to the CPU and says why.



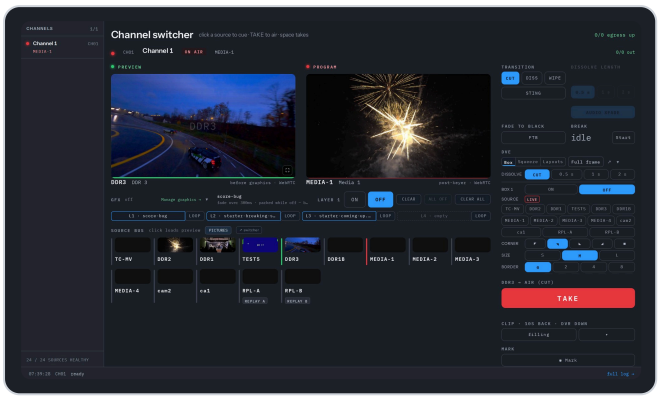
Inputs → + Add a feed. Name it, give it a tally code, then pick how it reaches you — SRT, RIST, RTMP, NDI, HLS, a browser stream, test bars or another feed's audio — and the form asks only what that way needs.

demo channel

06

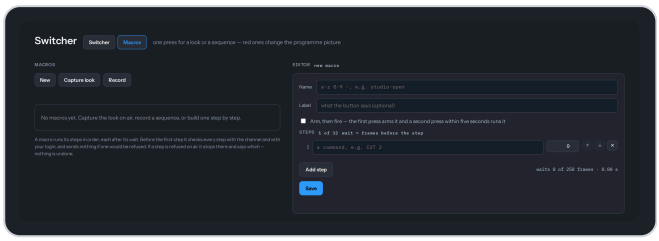
Switching, transitions & macros

A full production switcher with a preview/program workflow, transitions, DVE boxes and keyed scenes, and every take shown in advance.



Switcher. Preview and programme, the graphics layers, the source bus, and transitions and DVE with lengths in seconds. Box, Squeeze and Layouts sit beside TAKE.

example pictures



Switcher → Macros. Capture the look on air, record a sequence, or build it step by step.

demo channel

Preview, TAKE and CUT

Cue a source on preview and TAKE it with the armed transition, or CUT any source straight to air. Audio follows video, and a take lands at once.

Transitions

Cut, dissolve, SMPTE wipes and stingers: a clip plays over the take with the cut landing underneath, on a frame you choose. Stingers with transparency let air show through. Lengths are in seconds at the channel's own frame rate, and a new channel's dissolve is one second.

DVE, picture-in-picture and chroma key

Up to 8 boxes over the main picture with crop, border and a chroma key each; built-in layouts for interviews and L-bands, and your own saved as scenes. A box can dissolve from one source to another, and two boxes can swap places or sources as one move. A scene with an empty box is flagged before you take it.

Look before you take

A scene, a layout or a box move goes to preview first, drawn from the live pictures with its backdrop, boxes and artwork, and TAKE puts it up. The one direct action is a box's source, marked **LIVE**; on a box that is up it takes a second press.

A scene that goes with your take

ARM a scene and it lands with your next TAKE or CUT, on the same frame, with its main camera cued to preview. TAKE names the scene while it is armed. Schedules, playlists, failover and breaks never fire it.

Monitors that say what they show

If the programme monitor cannot show programme as it goes out, it shows the on-air source and says **BEFORE GRAPHICS** across the picture, so nobody reads it as air.

A second M/E bank

Switch the bus between **MAIN** and **M/E 2**: its own programme and preview over the same inputs, its own sound and graphics, with violet tally — for a regional opt-out or a second show.

Frame-accurate scheduling

Schedule a cut, a clip, a break, SCTE-35 or a layout for a time, and it fires on that frame. Every scheduled action is logged with its planned and actual time.

Macros

One press for a look or a sequence: capture what is on air, record steps, or build them one by one, with waits in frames. Every step is checked before the first one runs; red macros need a second press.

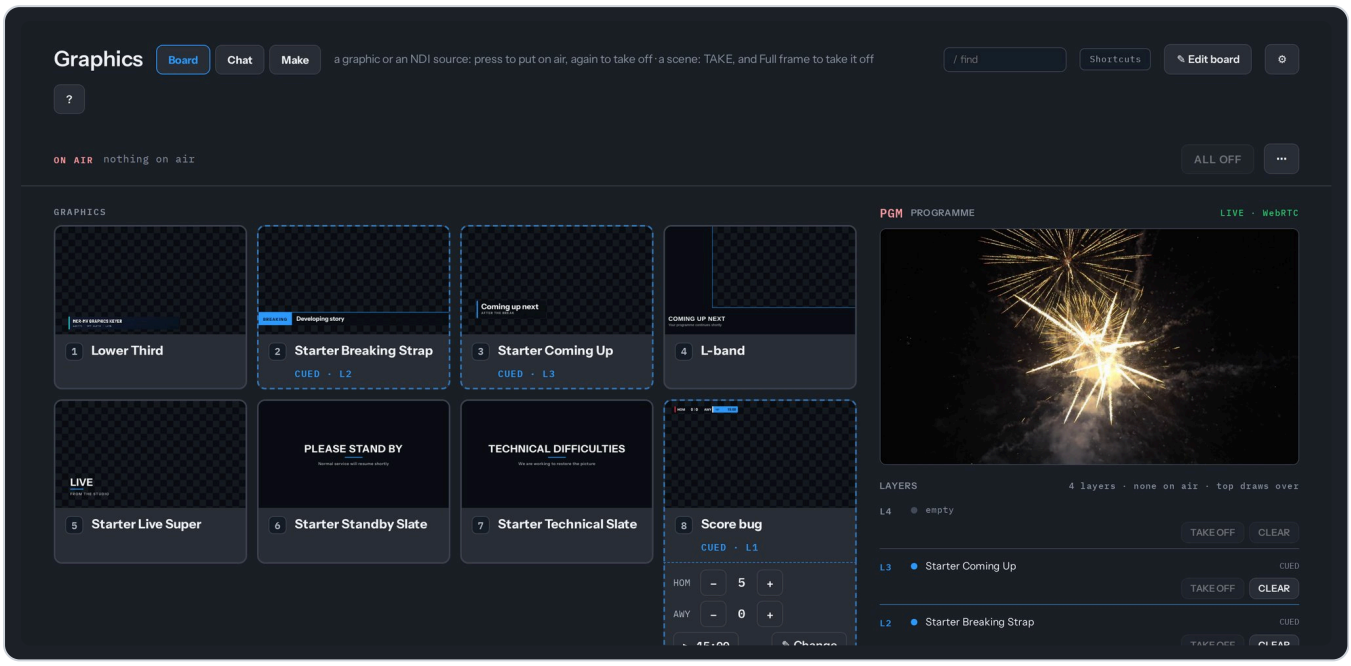
Fade to black, properly

FTB takes the graphics — bug, lower third, clock — down with the picture and puts back exactly what was up when it comes off. Captions stay on. The switcher always says why programme is black.

07

Graphics & DVE scenes

Graphics has two pages. **Make** builds them, and nothing on it goes to air. **Play** runs them, from a Board where each graphic is one press.



Graphics → Play. The Board: a button per graphic, named in words — lower thirds, straps, slates and score bugs with their own on-air controls — beside a live programme monitor and the layers, each showing what it holds. Press to air, press again to take off.

example pictures · credits at the back

Six keyed layers

Stills, clips with transparency, templates, web pages and NDI sources on six downstream layers. A new graphic on a busy layer fades the old one out first, on the Switcher's key bank as on the Board. ALL OFF fades everything and keeps it loaded.

Change words on air

✎ **Change** on a Board button edits a graphic while it is up. Score bugs have their own buttons — a cricket bug scores runs, wickets and balls in one press — and a score survives a re-take.

NDI graphics with alpha

A graphics PC's NDI output becomes a keyed layer, cut out by its own alpha channel, and reconnects by itself. NDI sources sit on the Board like any other graphic.

Build without touching air

Nothing on Make can reach air. An edit that would change something already on air, such as a data source a live button reads, names it and asks for a second press.

23 animated templates

Lower thirds (one name, two names, headshot, social), breaking strap, locator, now and next, countdown, title and quote cards, key points, poll results, league table, line-up, weather, credits, ticker, and football, tennis and cricket score bugs. Elements fade, slide, wipe or scale in and out.

Your brand, your fonts

Make → Brand holds the house brand — accent and plate colours and your own typefaces — and **Templates** builds your own, drawn by the same renderer that draws air. Built-ins are never changed.

Web pages on the Board

An HTML graphic — a Singular or Flowics overlay, a CasparCG-style template — can be a saved Board button with its own name and layer. The token in its address is hidden from every board view.

DVE scenes

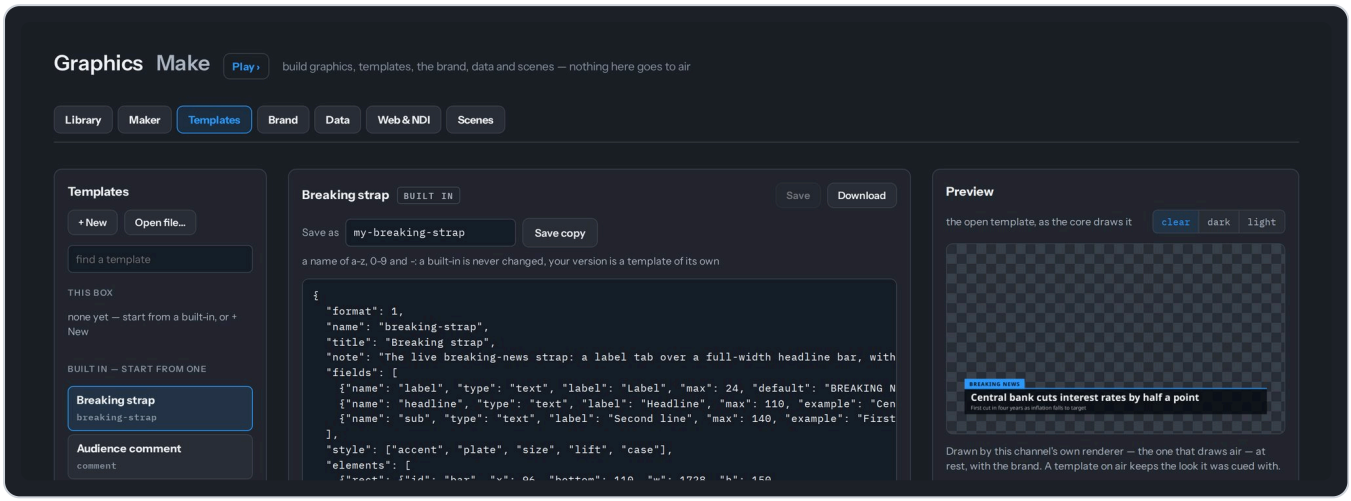
A scene is a whole look — picture-in-picture, side by side, up to nine pictures — over a backdrop and under artwork, built in **Make → Scenes**. On Play each scene is a tile with TAKE, CUE and ARM, and **Full frame** takes it back off. A take is one act: clips roll, keys set, boxes move and programme cuts together.

Live data and boards from a sheet

Link fields to a Google Sheet, CSV, JSON or RSS feed, or data your system pushes. A guest list in a sheet becomes one button per row, with names checked for fit at import, not on air.

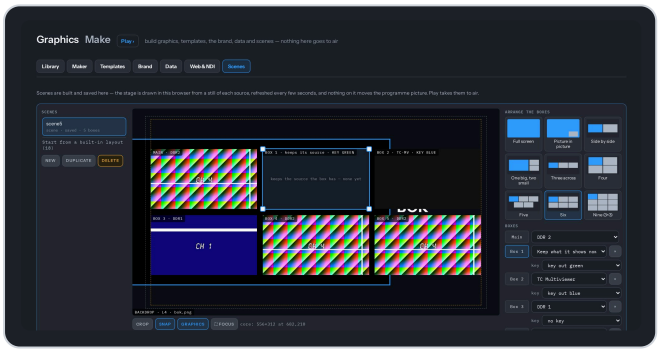
Audience comments on air

Comments from a YouTube stream or a Twitch channel arrive in a list. A producer queues one and TAKE puts it on a graphics layer; nothing reaches air by itself.



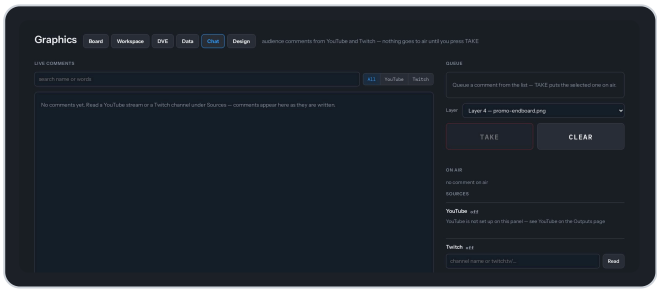
Graphics → Make. Library, Maker, Templates, Brand, Data, Web & NDI and Scenes. Here, Templates: start from a built-in or a blank template and preview it from the channel's own renderer.

demo channel



Make → Scenes. Arrange the boxes, choose each source and key, and drag on a stage drawn in the browser. Nothing here moves programme; Play takes the scene to air.

demo channel



Graphics → Chat. Comments from a YouTube stream or Twitch channel; queue one and TAKE puts it on a chosen layer.

demo channel

08

Audio desk, 5.1 & AC-3

A channel strip for every source under the switcher, programme processing to the delivery spec, buses for every feed that needs its own mix, and a 5.1 programme where the delivery needs one.



Audio desk. A strip per source — fader, mute, PFL/AFL, audio-follow-video, compressor, EQ, lip-sync delay and the aux sends — with monitor, talkback and master; below, the four aux buses and each dub language's bed, duck and voice levels. demo channel

Faders and mutes

A fader (−60 to +12 dB) and mute on every source. Meters and PFL are pre-fade, so a muted source still meters and you can tell "I muted it" from "it died". Strips with no feed fold away until asked for.

Loudness to spec

EBU R128 metering: momentary, short-term, integrated, loudness range and true peak, against −23 LUFS (EBU R128) or −24 LKFS (ATSC A/85). Correction and the loudness alarms are on from the start.

Dubs on the desk

Each dub language has a strip of its own and its bed, duck and voice levels; a routing grid decides which outputs carry which language.

AC-3, E-AC-3 and SMPTE 302M

AC-3 on any transport-stream output, signalled the DVB or the ATSC way; E-AC-3 and SMPTE 302M PCM on SRT and recordings.

EQ, dynamics and delay

3-band and parametric EQ, a compressor, optional noise suppression and a lip-sync delay (0–2 s, to the frame) on every source; EQ, compression and a true-peak limiter on programme.

Programme and aux buses

Programme with audio-follow-video, trim and master, plus aux buses fed from every strip for mix-minus, clean effects or commentary. Each bus can be listened to on its own.

5.1 programme

Run the programme bus in 5.1. Stereo sources are placed, never up-mixed into fake surrounds, and a standard Lo/Ro fold-down (ITU-R BS.775) feeds the stereo outputs.

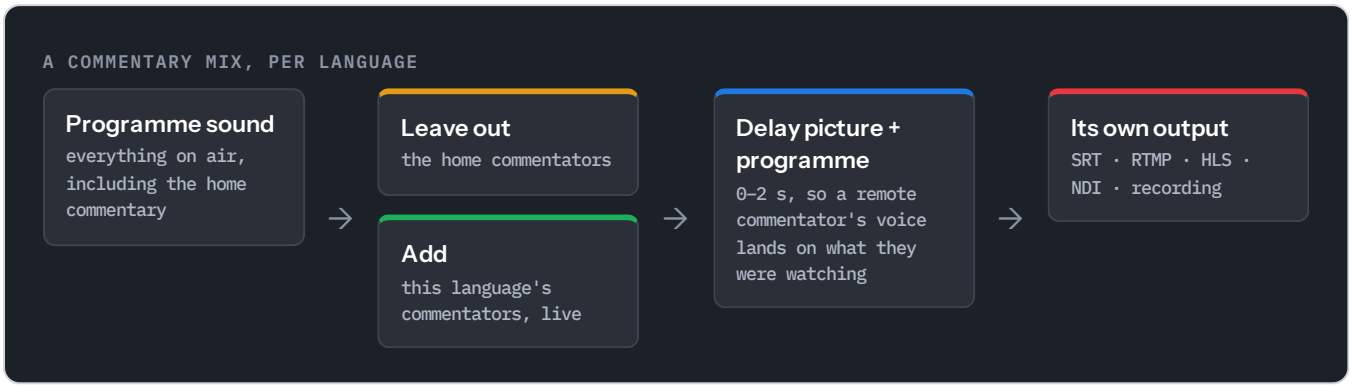
Multi-track outputs

Programme, each mix, each dub language, audio description and each aux bus can go to an output as separate tracks, with fixed PIDs and ISO 639 language labels.

09

Programme mixes, commentary & comms

One channel can feed several audiences, each with its own sound — and the crew can talk to each other and to guests without anything reaching air.



Follow mixes

A bus can follow programme with some sources left out and others added at their own level — your own commentary or interpreter mix — and any output can carry it.

International clean

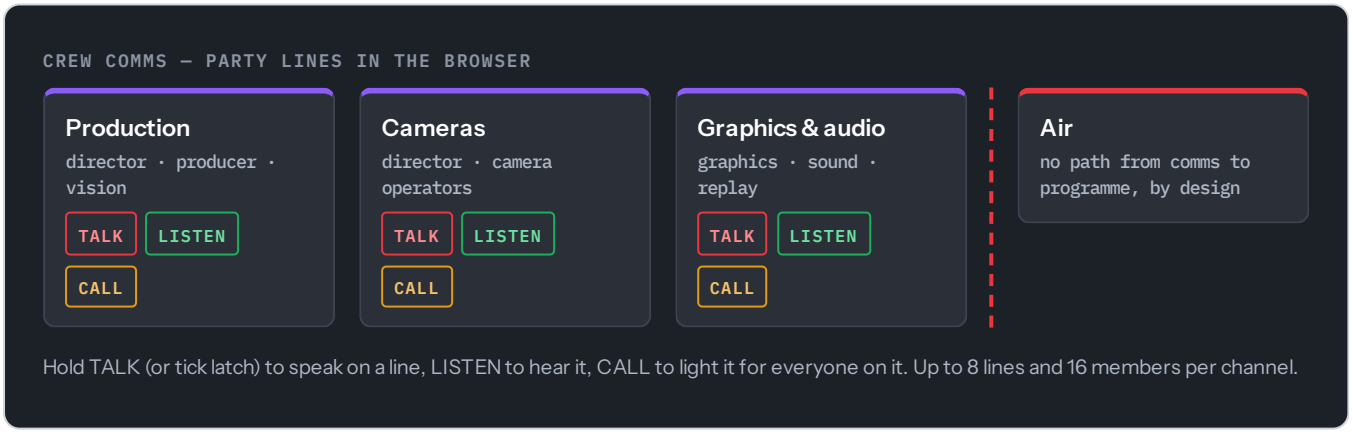
One choice on the new-output form fills in a clean picture with the music-and-effects bed, the feed rights holders ask for.

A mix in the player

Any bus can be offered to viewers in the browser player and as an alternate audio track on HLS OPTIONAL.

Mixes keep their labels

Each mix goes out as its own labelled track, with the ISO 639 language and a fixed PID, so an affiliate's decoder finds it every time.



Party lines OPTIONAL

Crew join from a browser with a headset; nothing to install. Lines are set per channel, and a member can sit on several at once.

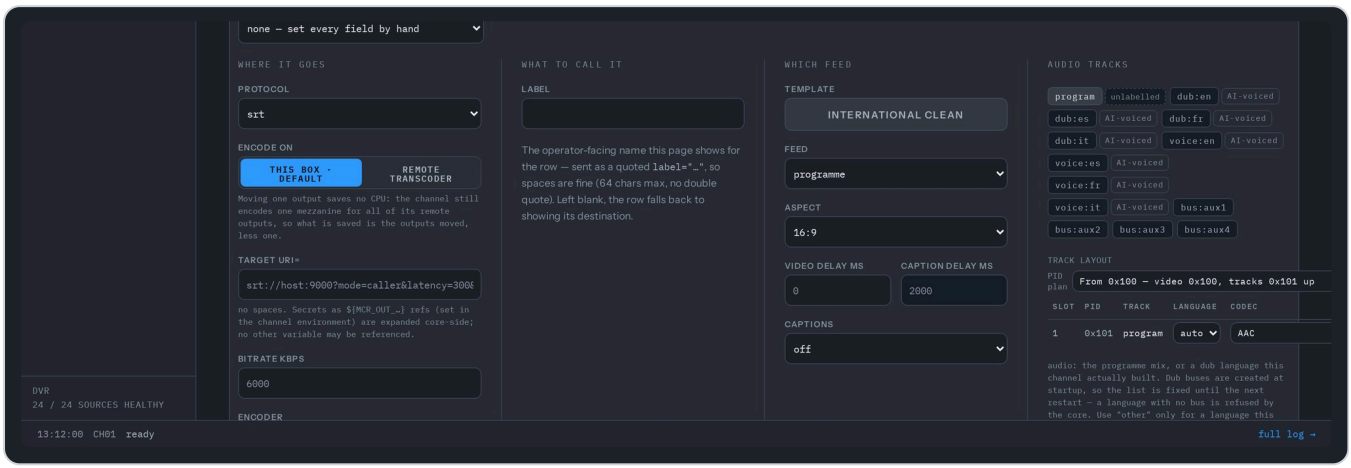
Talkback to guests

Hold a guest's button, or ALL, to talk to them. Only they hear you, their programme dips under your voice, and your microphone can never reach air.

10

Outputs, HLS & delivery presets

One programme, many destinations, running at the same time and added or dropped while on air. A failing output is contained and can never take the others off air.



Outputs → + New output. Where it goes and how (a destination preset, the protocol, which box encodes it), which feed and shape, and which audio tracks — programme, each dub and AI voice, each aux bus — each on its own PID with its language.

demo channel

SRT with strict CBR

Caller or listener, encrypted, multi-language MPEG-TS with fixed PID plans. Strict constant bitrate for headends and FAST platforms that require it.

HLS origin with ABR ladders

A live HLS origin with an adaptive-bitrate ladder, a subtitle track per caption language and an alternate audio track per dub. Segments sit on a UTC grid, identical on every origin.

Programme, clean or multiview

Any output can take the programme, the clean feed or the multiview wall, at 16:9 or as a 9:16, 1:1 or 4:5 crop that can follow the action, with captions embedded as 608 or 608+708.

Encode where it suits you

One switch for programme, outputs and monitors: Auto, CPU only, NVIDIA or Intel. CPU-only is a fully supported way to run. A failed hardware encode on an H.264 output falls back to the CPU and raises an alarm.

RTMP, RTMPS and Enhanced RTMP

Streaming platforms and CDNs, including HEVC and AV1 over Enhanced RTMP where the destination accepts it. Stream keys are masked everywhere.

Delivery presets

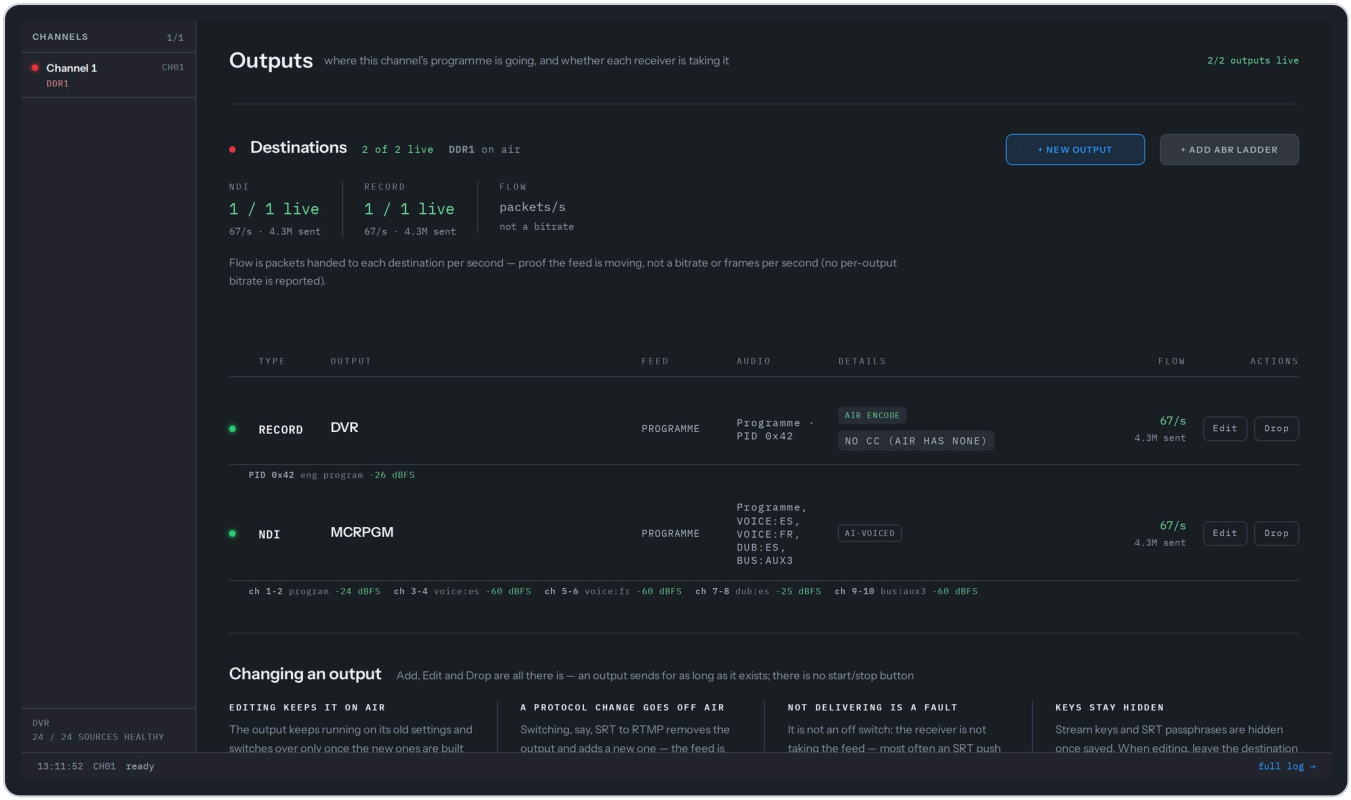
YouTube, Twitch and Facebook; Pluto TV, Amazon, Roku, Wurl and other FAST platforms; US cable and DVB headends; contribution HD. Each fills in its destination's published spec and warns about anything that does not match.

NDI, near or far

Programme or any bus back onto the network as NDI — from the channel box, or from a transcode node on another site's network BETA.

MPEG-2 and interlaced legs

An output can be MPEG-2 or 1080i for a headend that still needs it, alongside progressive H.264 and HEVC from the same programme.

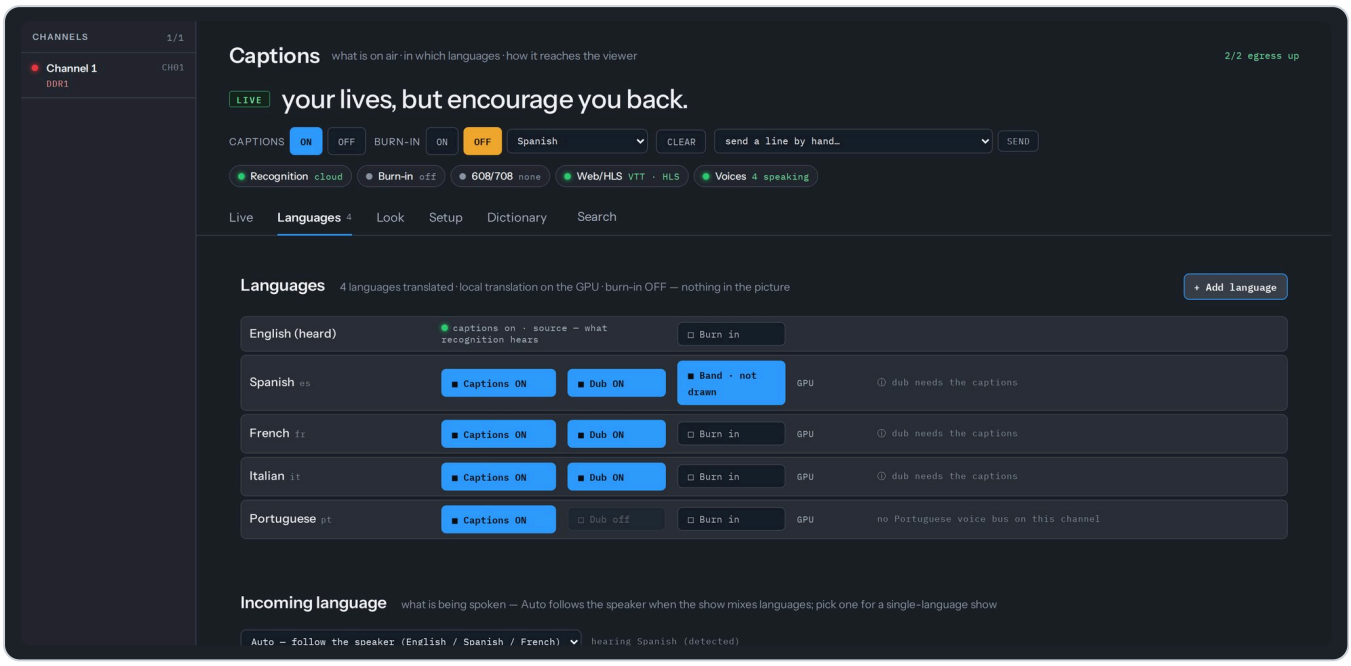


Outputs. Each destination with its feed, its audio tracks and PIDs, what it carries (AI-voiced tracks, captions) and whether the receiver is taking it. demo channel

11

Captions & translation

Live captions and translated subtitles, built into the channel rather than bolted on — and they follow the speaker when the language changes.



Captions → Languages. Every language with its captions, dub and burn-in in one row each, and the incoming language set to **Auto — follow the speaker**: here it hears Spanish and the English captions carry the translation.

demo channel

Live captions

Programme audio is transcribed live and embedded as CEA-608 and CEA-708 — in each output, or in the air encode itself — written to WebVTT and HLS subtitle tracks, or burnt into the picture. Captions already in a source can be passed through.

Follow the speaker BETA

Set the incoming language to Auto and a guest who speaks Spanish on an English channel gets English captions translated from the Spanish, and the Spanish dub carries their own voice. A change of language is caught within a few words, so captions and voices follow it live. Or name the language for a single-language show.

28 languages, told straight

The catalogue lists 28 languages for live recognition; accuracy has been measured for English and Spanish so far. From English, local translation covers Spanish, French, German, Italian, Portuguese, Mandarin and Arabic, and local voices cover English, Spanish, French, Italian, Portuguese, Hindi, Japanese and Mandarin. So a fully local dub from English works in five languages: Spanish, French, Italian, Portuguese and Mandarin. Others use cloud services. Languages CEA-608 cannot carry go to subtitle tracks and burn-in only.

Burn in any language

The burnt-in band can show the speaker's words, a translation instead, or both as a bilingual pair. One pick from the BURN-IN menu starts the translation and puts it on screen.

Translation in whole sentences

Set the targets once. Each language becomes a caption track, a subtitle rendition, a burn-in choice and, if you want, a dub. Translations are made from whole sentences, so they read like sentences.

Where captions are going

One card on Captions → Live shows recognition, burn-in, 608/708 per output, the player, HLS subtitles, recordings and each dub — and says plainly if nobody is getting captions.

Captions in the player

The browser player shows captions on by default, big and clear of graphics, with the speaker's own language arriving word by word.

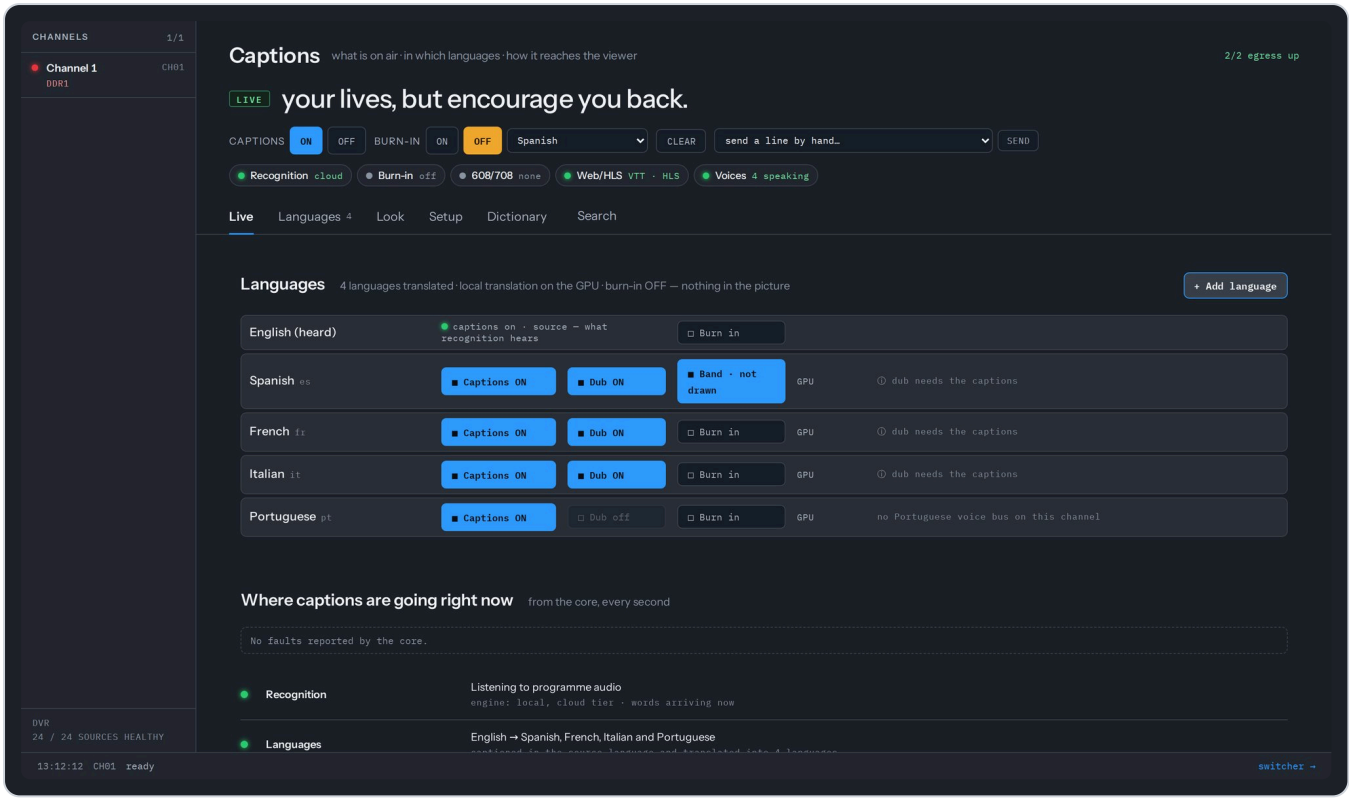
A dictionary that is yours

Captions → Dictionary teaches the recogniser names, places and teams, and fixes the words it keeps getting wrong. Switch dictionaries live between shows.

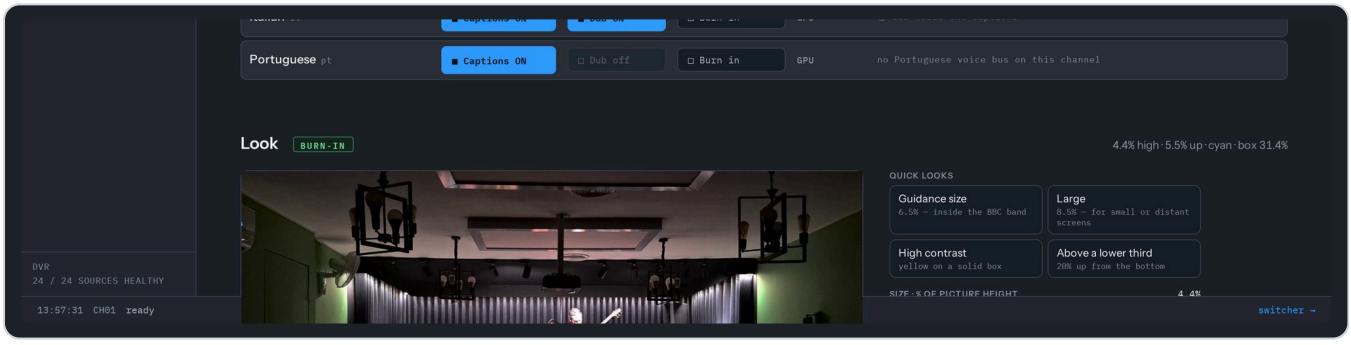
> **Captions never take the programme down.** If a speech or translation provider fails, captions pause and an alarm says so, while the picture and sound on air carry on untouched.



Live captions on the programme. Words recognised as they are spoken and shown on the player, while the same captions go out embedded in the outputs. Footage: NASA (public domain).



Captions → Live. The line on air, the BURN-IN switch with its language, the status of every path, and the languages strip. demo channel



Captions → Look. How burnt-in captions look on the live picture — size, position, colours and box — with quick looks such as guidance size, large and high contrast. demo channel

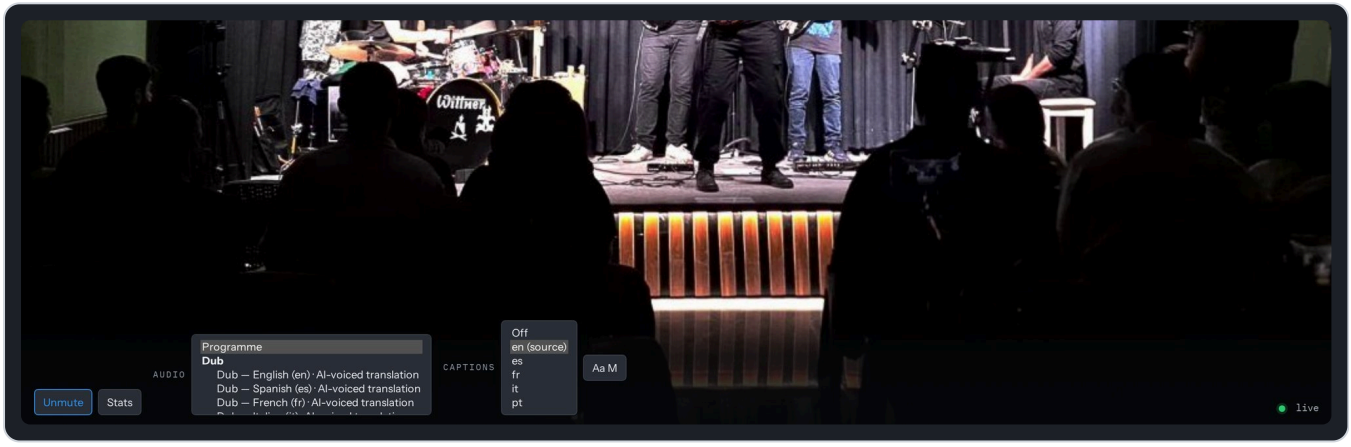
12

AI voices & audio description

A spoken track in every language the channel translates, mixed over the programme sound so the room and the music stay in — and labelled as AI wherever it appears.

ONE PROGRAMME, EVERY LANGUAGE — AN EXAMPLE CHANNEL

Programme	EN the main mix, on every output
Captions	EN live, embedded as 608/708 · EN ES FR IT WebVTT and HLS subtitles the viewer switches on
Voices	ES FR IT AI-voiced tracks on SRT, HLS alternate audio, the player and recordings
Audio description	EN programme with spoken description, its own labelled track



Player. The viewer picks the sound — programme, or an AI-voiced dub in English, Spanish, French or Italian — and the captions: off, the speaker's own language, or a translation. Changing the sound keeps the picture playing.

demo channel · lists opened for the picture

A voice per language

Each translation language can have its own synthesised voice, carried as its own audio track. Each language stays off, and costs nothing, until you switch it on. Each speaker can have a voice of their own.

Mixed like an interpreter

The programme bed dips under the voice and comes back between lines, with depth, attack and release set per language on the Captions page or the desk. A line that arrives late is spoken late rather than dropped, and an alarm says if a language starts falling behind.

Health you can see

A voice that is not running is greyed out with the reason — "no French voice running" — rather than a silent track going to air.

Local or cloud voices

Local voices (Kokoro) run on an NVIDIA graphics card NVIDIA GPU; cloud voices from ElevenLabs, OpenAI or Mistral need no card at all.

On the desk

Each dub has a strip on the audio desk and can be sent to an aux bus; a routing grid decides which outputs carry which language.

AI voices are labelled

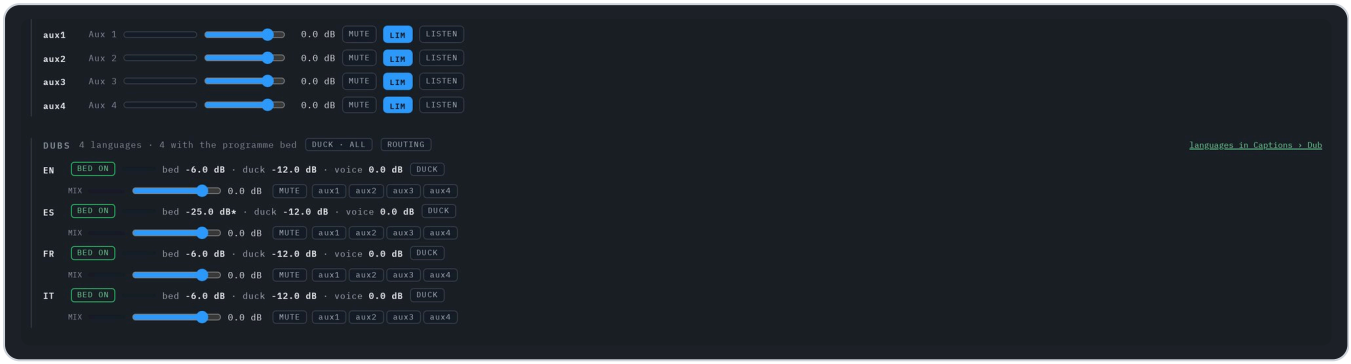
Every synthesised voice is marked as AI-voiced wherever it appears — player, outputs, captions page and guide — in your channel's own wording, with an optional spoken disclosure.

Audio description

An AD service per language, voiced by an AI describer that watches the picture or by lines an operator types, with DVB audio-description signalling. Every line is logged for 30 days.

Who is speaking

Captions mark speaker changes with >>, a colour per speaker or the speaker's name — on an AI node `NVIDIA GPU` or with Speechmatics. Presenters are enrolled only with their recorded consent.



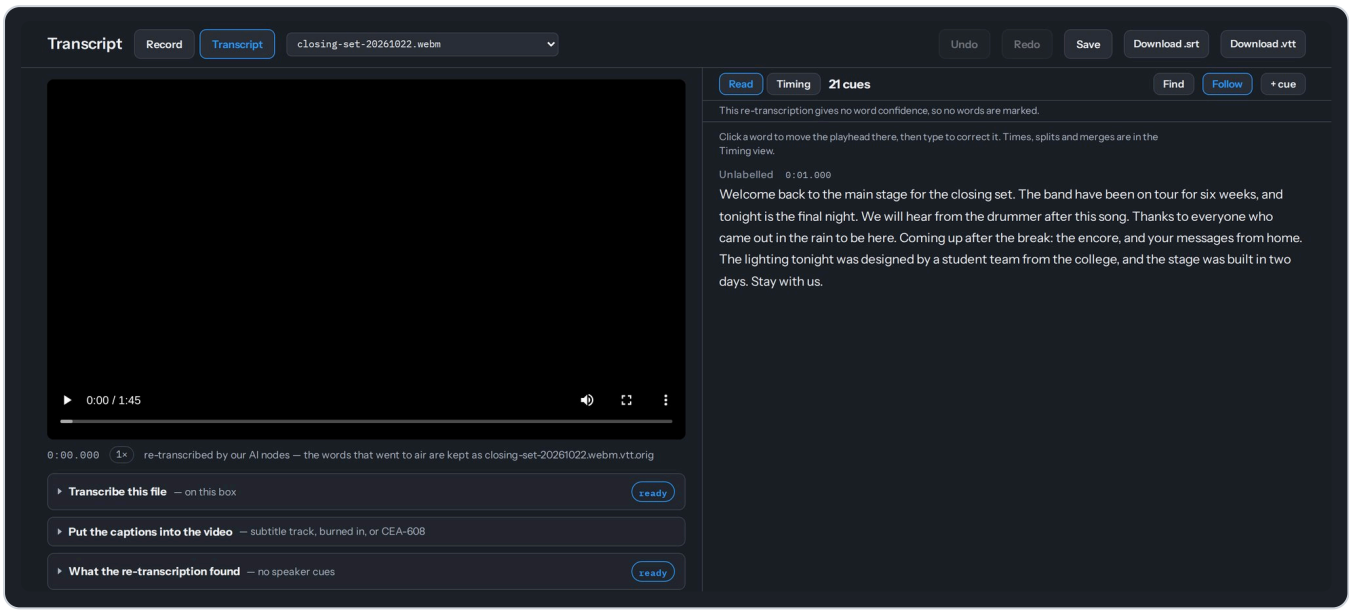
Audio desk → Buses and Dubs. The aux buses, then each dub language with its bed level, duck depth and voice level, its mix fader and the buses it feeds.

demo channel

13

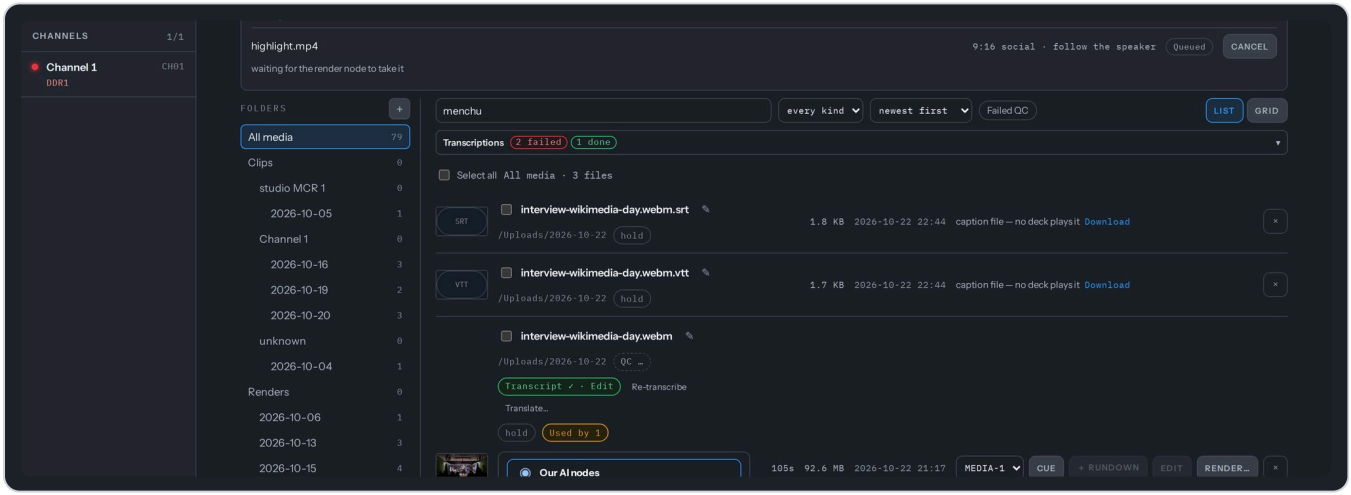
Transcripts, the transcript editor & search

What was said is kept, corrected, translated and searchable — for recordings, clips, media-library files and the feeds coming in.



Transcript editor. The recording beside its words in the Read view: click a word to jump there and type to correct it; Timing shows the cues. Below, transcribe the file again, put the captions back into the video, and the speakers found.

example picture and text



Media → Library. A file with its finished transcript (and the .srt and .vtt beside it), and **Transcribe** asking who should do it: the AI nodes, or a cloud provider whose cost is shown before anything is sent.

demo channel

An editor for words

A Read view in paragraphs per speaker and a Timing view for the cues; find and replace; rename or merge speakers; undo; low-confidence words underlined to check; and a clip made from the words you select.

A transcript beside every clip

A clip starts with the captions that went to air during it. The words that aired are always kept, so a corrected copy never replaces the compliance record.

Transcribe any file

Transcribe a video in the media library on an AI node `NVIDIA GPU` or through a cloud provider. Watch it progress, cancel it, and resume it after an interruption without starting again.

Translate a transcript `NVIDIA GPU`

Translate a finished English transcript into the channel's other languages, as subtitle files ready to burn or deliver.

Speakers on long files NVIDIA GPU

Speaker labels on files up to four hours long, found on the AI node's graphics card.

Summaries and chapters NVIDIA GPU

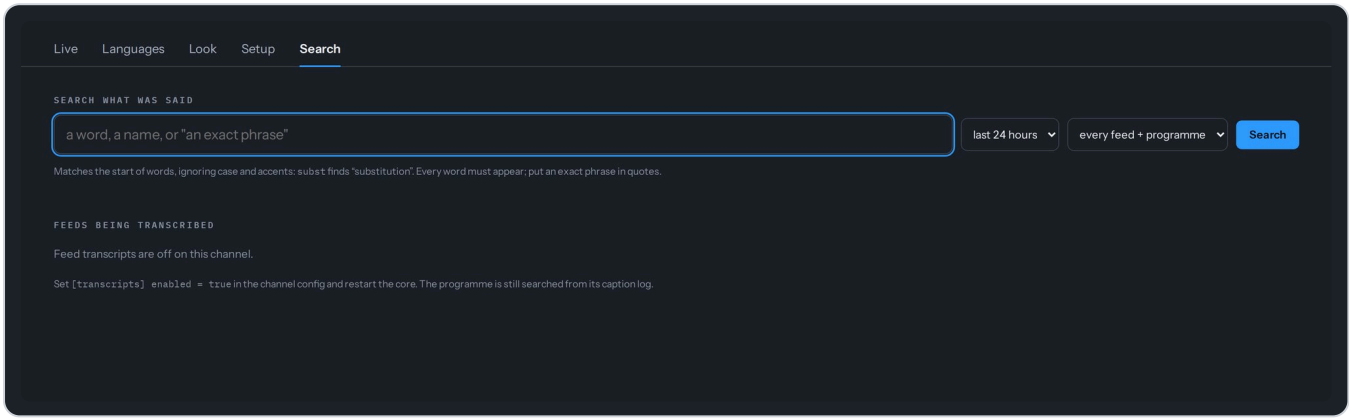
A summary, topics and chapter marks from a local language model on the AI node, so nothing leaves your network.

Correct, then re-burn

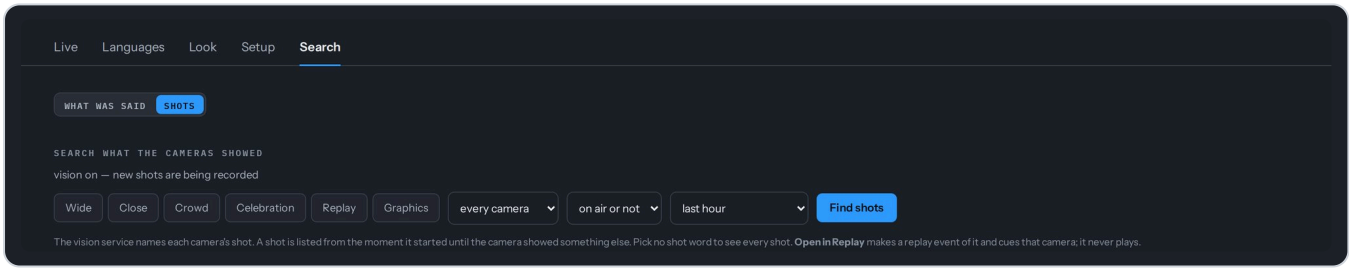
Fix a name once and add it to the dictionary, so the recogniser gets it right next time. Make a new file with the corrected captions; the original recording is never changed.

Search what was said

Captions → Search finds a word, a name or an exact phrase in the feeds, the programme's caption log and library transcripts, over the last hours or days, and jumps straight to the moment.



Captions → Search. A word, a name or "an exact phrase", over a time window, across every feed and programme. demo channel

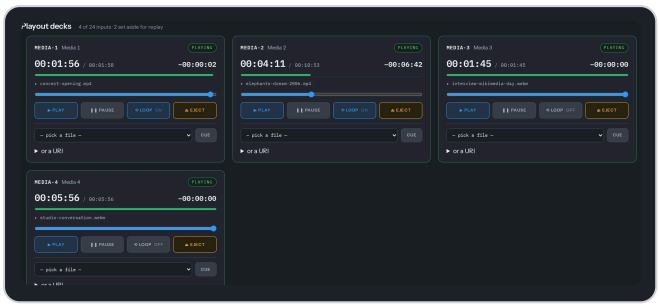


Search → Shots NVIDIA GPU. Search what the cameras showed: wide, close, crowd, celebration, replay or graphics, by camera, on air or not, over a time window. **Open in Replay** makes a replay event and cues that camera; it never plays. demo channel

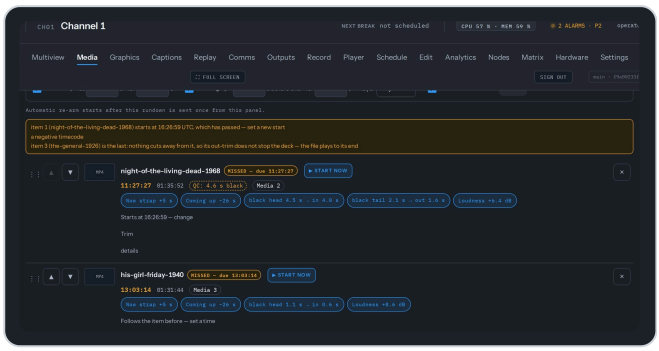
14

Playout, rundown & scheduling

Run a linear channel: packages on playout decks, a rundown built from the library, breaks that signal downstream, and actions scheduled to the frame.



Playout decks, popped out. Every deck with its file, elapsed and remaining time and a progress bar; cue, play, pause, loop and eject from a window of their own.



Media → Rundown. Each item with its start time, deck, the black head and tail found by QC, its loudness correction and its straps.

public-domain films

A rundown from the library

Add files from the library, drag them into order, trim in and out, and see the wall-clock start of each. Decks alternate by themselves, and "Now" and "Coming up" straps can ride on each item.

A rundown that is safe to run

A NOT ARMED banner until it is armed; after a restart it re-arms safely; a late start asks whether to start now, pick a time or cancel; a file already on air by hand is adopted, not restarted.

Ad breaks and SCTE-35

Start a break and the channel plays bumper-in, break and bumper-out, counts down and returns. Splice insert and time signals with segmentation descriptors go on the outputs that carry them.

Slate backstop OPTIONAL

Arm it in Settings → Air safety and, if everything feeding the channel is lost, it goes to a slate rather than to black, judged on the source's own frames.

Air-ready files

Each file's loudness is normalised as it plays, black-and-silent heads and tails are trimmed, and chips warn of a file that does not match the channel — "25 fps on a 59.94 channel", interlaced, 4:3.

Decks that handle any file

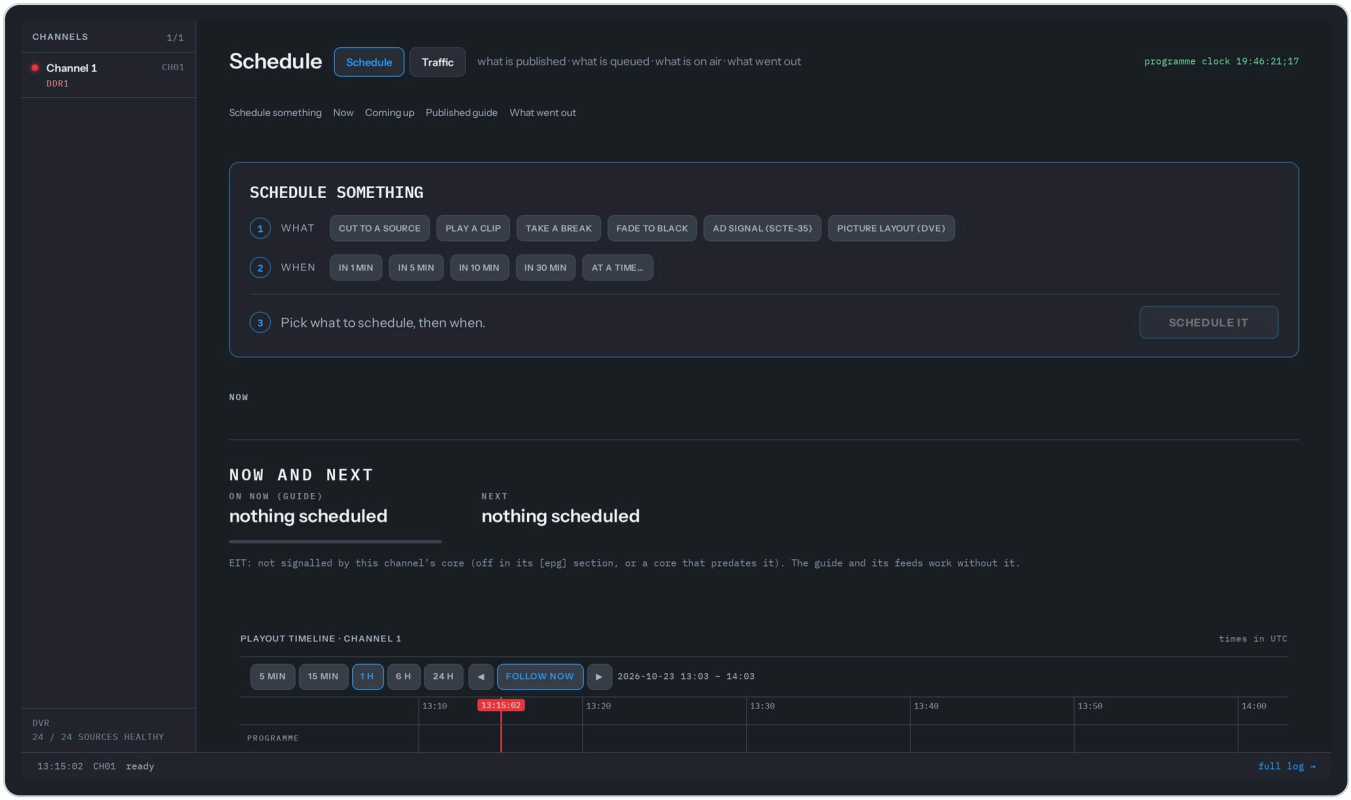
Interlaced, non-square-pixel and phone-rotated files are deinterlaced, set upright and pillarboxed or letterboxed rather than stretched.

Schedule something

What (a cut, a clip, a break, fade to black, an ad signal or a layout) and when (in a minute, or at a time) — two choices, and it is on the timeline, locked to the wall clock.

Now and next

The schedule shows what is on now and next from the programme guide, or from the rundown when there is no guide.

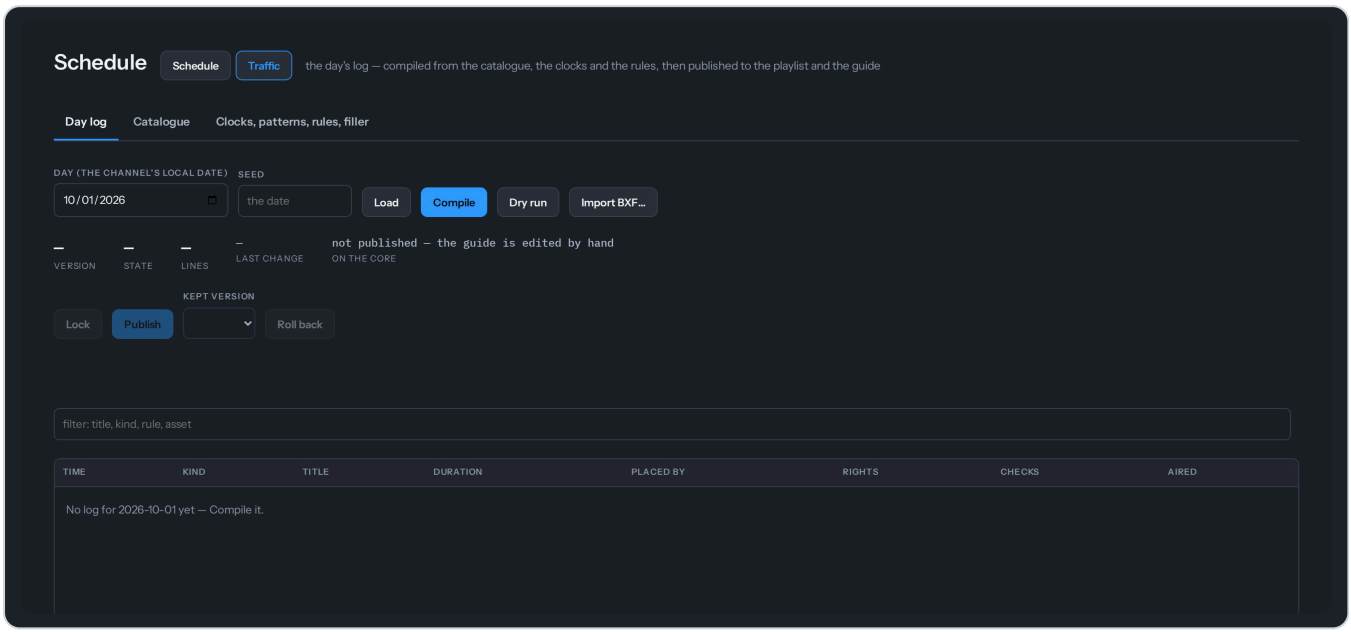


Schedule. Schedule something in two choices — what, then when — with now and next and the playout timeline below. demo channel

15

Traffic, programme guide & FAST delivery

Build the day's log from a catalogue and rules, publish it to the playlist and the guide, and deliver the guide in the forms platforms ask for.



Schedule → Traffic. The day log compiled from the catalogue, clocks, patterns, rules and filler — dry-run it, publish it, lock it, roll it back, or import BXF. demo channel

A traffic layer

Hour clocks, ad-break patterns, placement rules, filler and rights windows compile a catalogue into a day log. Every line says what placed it and whether it passed its checks.

Publish, lock and roll back

A compiled log is published to the playlist and the guide together. Versions are kept; a published day can be locked, and the last version is one press away.

Guide safety

The next 48 hours are frozen: changing them needs an admin or an engineer and a reason. Warnings flag gaps, short blocks and missing ratings.

House-format conform

Traffic and the rundown play a file's house-format version when one has been made, so off-format material airs at the channel's own rate and shape.

Clean heads and tails

A film's black-and-silent head and tail are trimmed automatically, and a spot or promo with one is flagged for traffic to fix rather than cut, so the break still fills to the second.

BXF import

Take the schedule from an existing traffic system in BXF, alongside CSV and XMLTV guide imports. One bad row refuses the whole file, with its line number.

Guide feeds and exports

XMLTV and Roku feeds behind their own revocable tokens, and EPG exports in the forms Samsung TV Plus and Pluto TV take; now-and-next EIT inside the transport stream for set-top boxes.

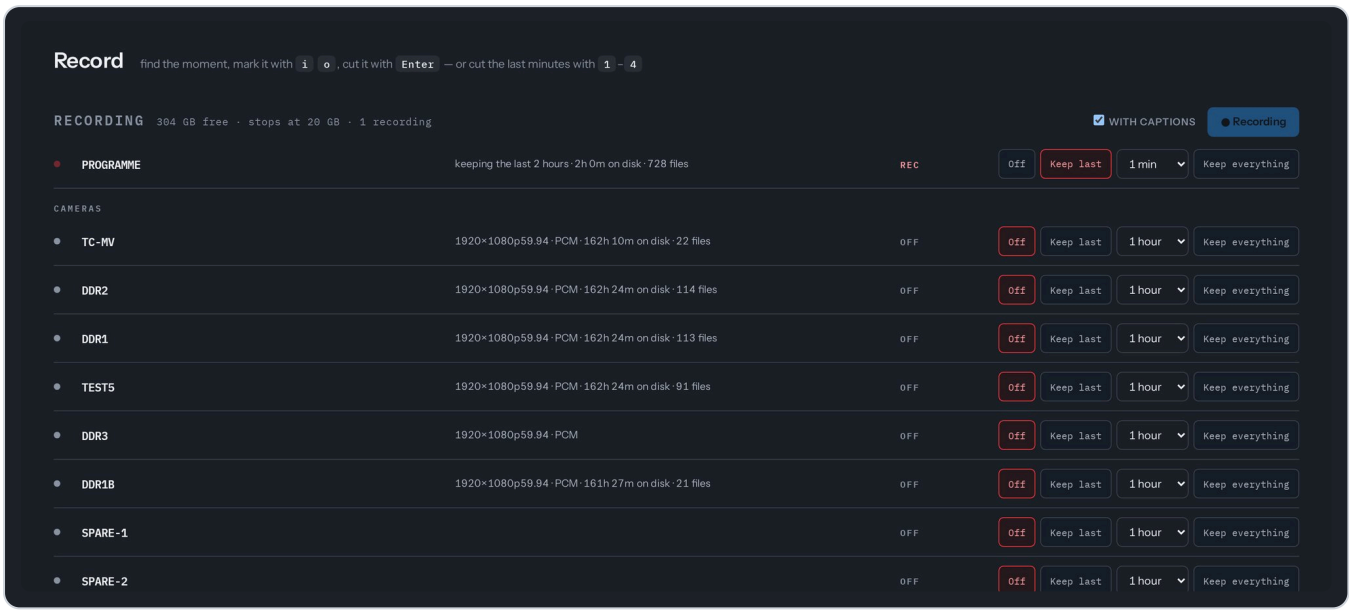
AI drafts, people publish

An AI agent with the right token can prepare a traffic log for review, but can never lock, publish or roll one back.

16

Recording, clips & highlights

Record programme and every camera, cut clips while the show is still on air, and let the channel propose the moments worth keeping.



Record. Programme and per-camera ISO recording, each kept for a window you choose, with captions and the disk's headroom shown. demo channel

Programme recording

Keep a rolling window or an archive (90 days unless you choose otherwise). Recordings carry the same tracks, language labels and captions as the feed they record. A DVR can record the air encode itself, with no second encode.

Air comes first

Recording stops at a disk floor you set, and a recorder that cannot start or a stalled disk never holds the channel off air.

Mark in, mark out

Browse the recording with J / K / L, step frame by frame, mark I and O. A clip is cut on the channel at the nearest keyframe in seconds; a render node cuts it on the exact frames.

Highlights with a score

Moments are proposed and ranked critical, major or minor from the signals you choose — loudness, keywords, crowd sound, a score change, a replay roll, picture motion and shot type `NVIDIA GPU`. Each says why in words: “Loudness · clear”, “Said: goal”, “Celebration shot”. Detection is on for a new install, and the director always decides.

ISO recording

Record any camera on its own — SRT, NDI, HLS, WebRTC or a guest — before it is switched, and arm or disarm it during the show.

Quick clips

Cut the last 15 s, 30 s, 1 min or 2 min with one key, from the Record page, the switcher or a Stream Deck. Clips from several angles of the same moment are one more press.

Captions on clips

A clip keeps the captions that were on air during it, with a subtitle file per caption language, embedded 608 or burnt into the picture.

One page for highlights

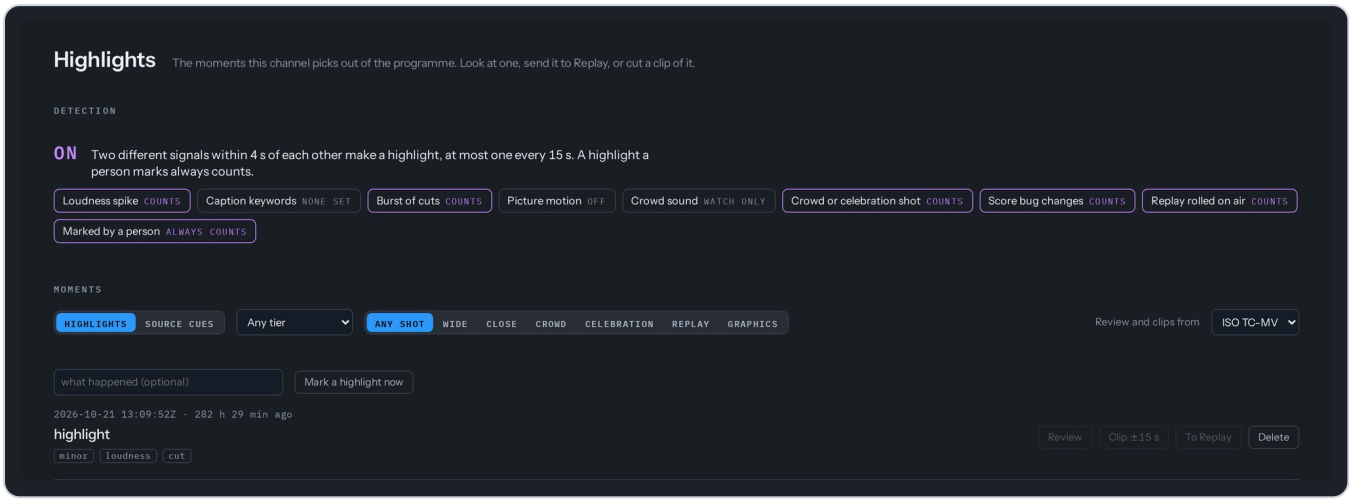
The Highlights tab says what is switched on and what each signal does now, lists every moment with tier, shot and on-air filters, and holds the rules.

Pre-cut for social BETA

A top highlight can cut itself and render 9:16 and 1:1 versions on a render node, ready to check and publish. Nothing goes to air or to social by itself.

From a moment to air

Review plays 15 seconds either side on the page, **Clip ±15 s** cuts it, and **To Replay** makes it a replay event to cue or add to a reel. Each says how far back its recording reaches.

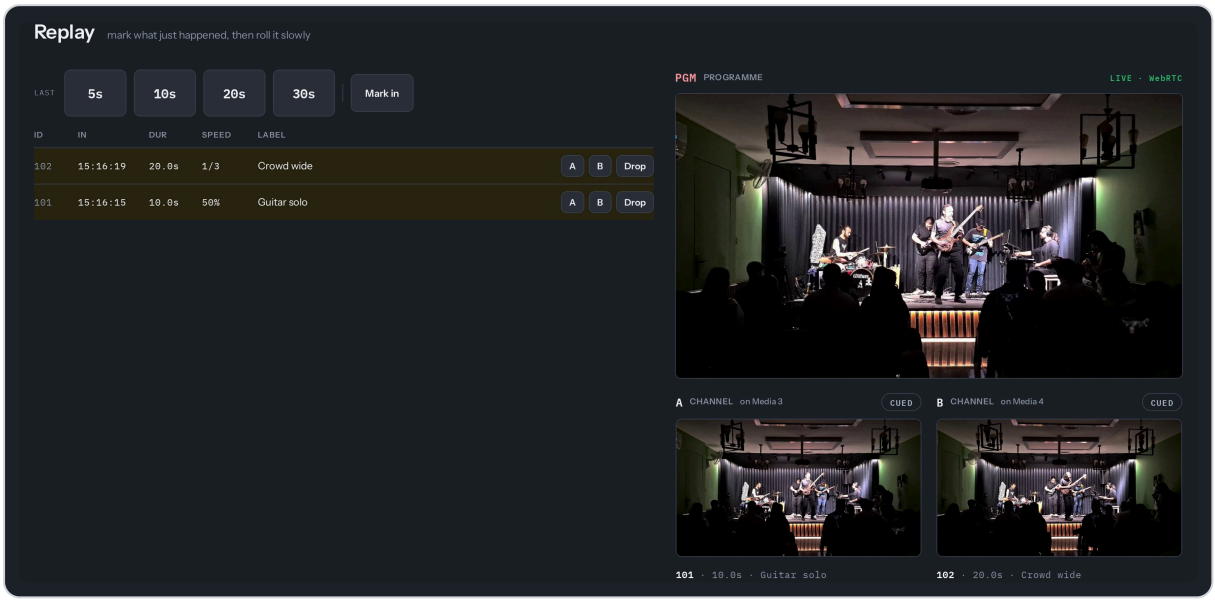


Highlights. Detection first: on or off, and what each signal does now. Then the moments, with tier and shot filters. Review, Clip and To Replay are grey here because the demo channel is not recording. demo channel

17

Instant replay

Mark what just happened and roll it back slowly, in the way replay operators already work. Replays come from the channel's rolling recording, so how far back you can reach is set by how much it keeps.



Replay. Two marked moments in the event list, and both replay channels cued: A at half speed, B at a third, beside programme.

example pictures · credits at the back

One press, after you have seen it

Mark the last 5, 10, 20 or 30 seconds — the biggest buttons on the page, because you mark once the moment has happened. Mark in and mark out are there for exact ends.

Any speed

Set the speed as a percentage or an exact fraction — 50%, 1/3, 1/4 — or 0 to hold a frame, live while it rolls. Play two channels together, or follow another angle mid-roll.

Trim before air

Click an event to scrub it, set a new in and out, and play from there, with J / L and I / O on the keyboard or a jog wheel.

Jog and shuttle

Plug a Contour ShuttlePRO or ShuttleXpress into the browser, or a MIDI surface, and teach it new controls.

Replay ready BETA NVIDIA GPU

When the detector finds a moment — a crowd or celebration shot with the roar or the score — it makes the replay event and puts up a card: “Celebration — CAM 3 · Ready on A”. If replay A is idle it is cued from the best angle; if not, the card says “Replay busy”. It never plays: PLAY A is the operator's press. Set it to suggest only, or off.

Every angle on the same frame

Cue programme or any camera, and every angle lands on the same frame, with a per-camera offset to line up a late feed. A strip of stills shows each angle at the marked moment.

True slow motion

A 100 or 120 fps camera, recorded on its own, replays at native speed: one picture on air for every frame the camera took.

Highlight reels

Add events to a reel and play them back to back, cut or dissolved, with Play reel, Next and Stop.

To air like any source

A replay is cued on preview and taken, so it gets the armed transition, tally, the as-run log and the fault slate. Sound is muted below real time.

Fed by highlights and shots

A highlight's **To Replay** and a shot's **Open in Replay** arrive as ordinary events, ready to trim, cue or add to a reel, and a still can show any moment the recording still holds.

18

Editing, render nodes & publishing

Turn marked moments into a finished clip, in the shape each platform wants, get every file into house format, and send it where it needs to go without leaving the panel.

A cloud editor

Stitch several moments into one clip: reorder, trim, split and undo, with cut, dissolve and dip joins, audio lanes with levels and delivery loudness, speed ramps, and keyframed pan and zoom.

Rendered off the air box

Every render — edits, captioned clips, exports, social versions — goes to a render node, GPU or CPU, which checks each output before handing it back. With no node connected, the channel box renders at low priority, and the panel says where each job runs.

Social and house versions

9:16 vertical and 1:1 square that follow the speaker's face, a house master in H.264, HEVC or AV1, and DNxHR and AVC-Intra 100 masters for post-production.

Publish to storage

Send a clip, its poster and its caption files to S3-compatible storage. Uploads resume and retry by themselves.

From markers, recordings or files

Start an edit from ticked markers, a marked range of the recording, an exact span from a log sheet, or any video in the media library.

House format, checked and fixed

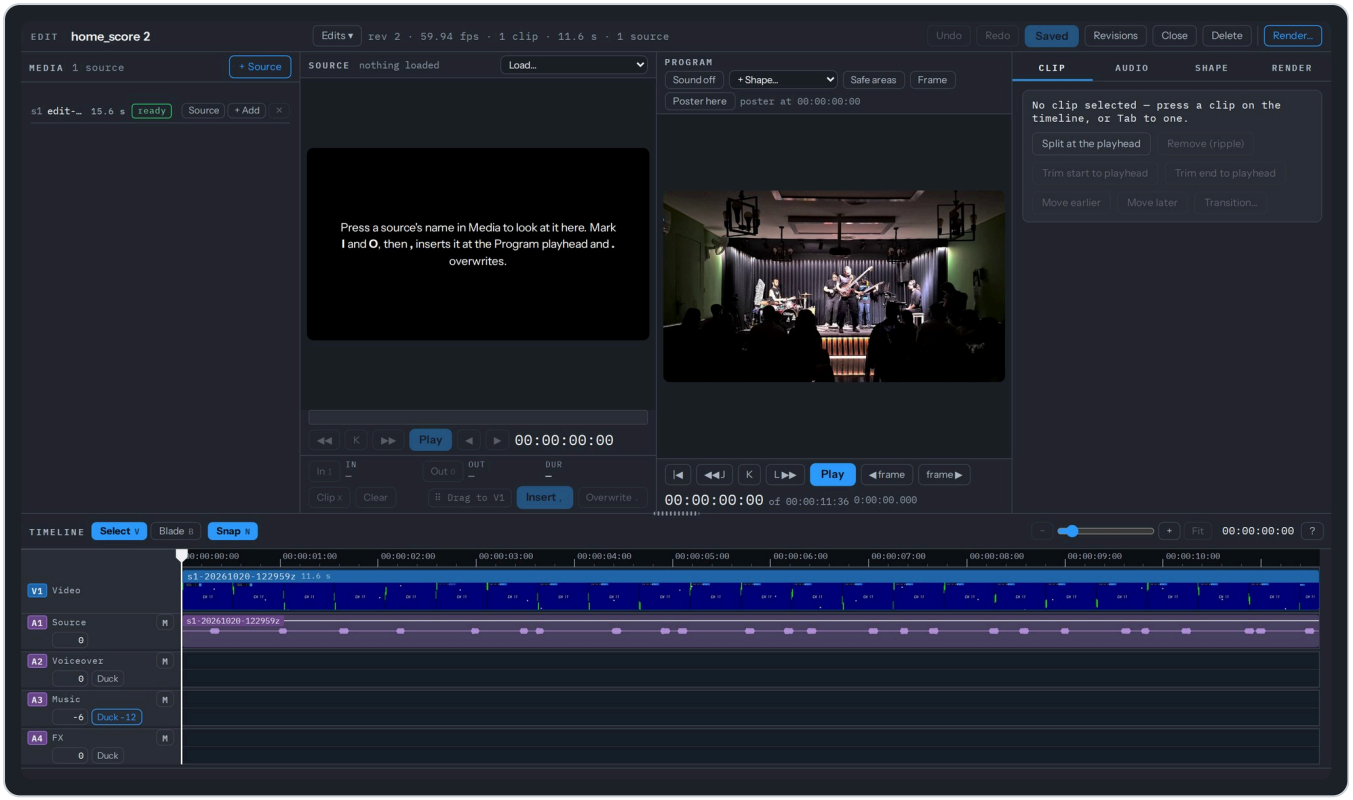
Every library file is judged against the channel's house format — frame rate, interlace, pixel shape, rotation, variable frame rate, codec, size — and **Conform to house** makes a matching copy on a render node, checked before it is used.

A library that is looked after

Folders, bins, tags and a trash; retention rules, legal hold and a cold-storage archive; and a quality check on every file for loudness, black and silent ends, and format. Any file can be downloaded from its row.

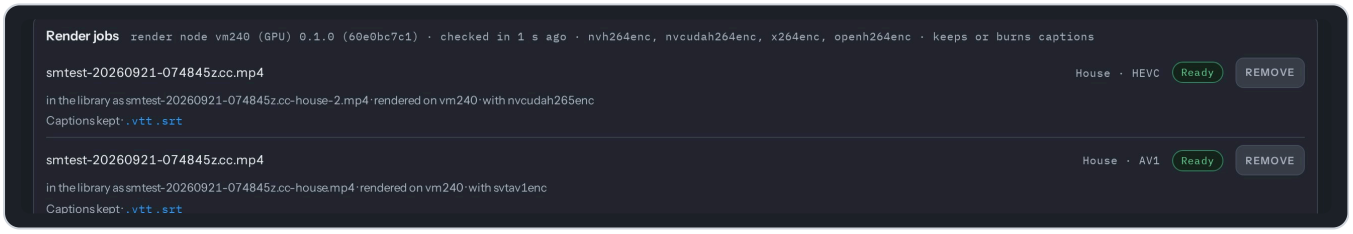
Publish to social BETA

YouTube, TikTok, X, LinkedIn, Facebook, Instagram and Threads through your own developer apps, with titles, captions and the AI-content flag. Not yet proven against every live platform; some need the platform's own approval first.



Edit. Source and programme monitors, the clip, audio, shape and render tools, and a timeline with video and four audio lanes — source, voice-over, music ducked under speech, and effects.

demo channel



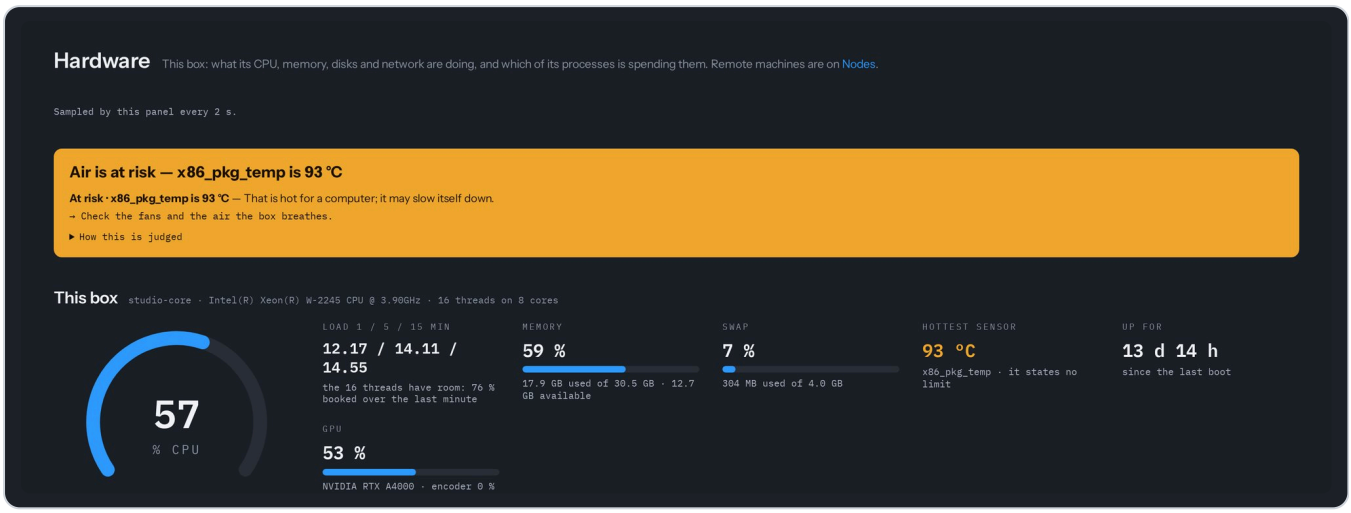
Media → Render jobs. A GPU render node making HEVC and AV1 house versions, with the captions kept as .vtt and .srt.

demo box

19

Alarms, hardware & operations

Every fault is reported in plain English, every alarm is yours to tune, and the box tells you what it is doing — and whether air is safe.



Hardware. A verdict in words first — here a warning that the CPU is running hot — then CPU, memory, graphics card, the hottest sensor and every thread, with the channel's CPU split by job. demo box

A verdict, not a dashboard

Fine, running tight, air at risk or needs attention — judged from fixed limits and said in words, with how many threads are to spare and what to do about it. The channel's CPU is split by job — monitors, air, multiview, NDI, inputs, outputs, audio, captions, graphics — so you know what to turn down.

Alarm history

Every raise and return is kept on disk (90 days by default), so a fault last night can still be read this morning.

As-run and audit trail

Every on-air action is logged with who or what did it, including a Stream Deck or a token. A separate audit trail records sign-ins, settings and administration.

Metrics, webhooks and syslog

Prometheus metrics, signed alarm webhooks, RFC 5424 syslog to your SIEM, and an independent monitor with a dead-man's switch.

Alarm bus and policy

Prioritised alarms (P1 is the worst) in plain English. A flapping source is one alarm with a count. ACK, CLEAR and SNOOZE. More than fifty checks, each yours to turn on, re-prioritise and time; a new channel starts with the air-dark and loudness checks on.

Analytics

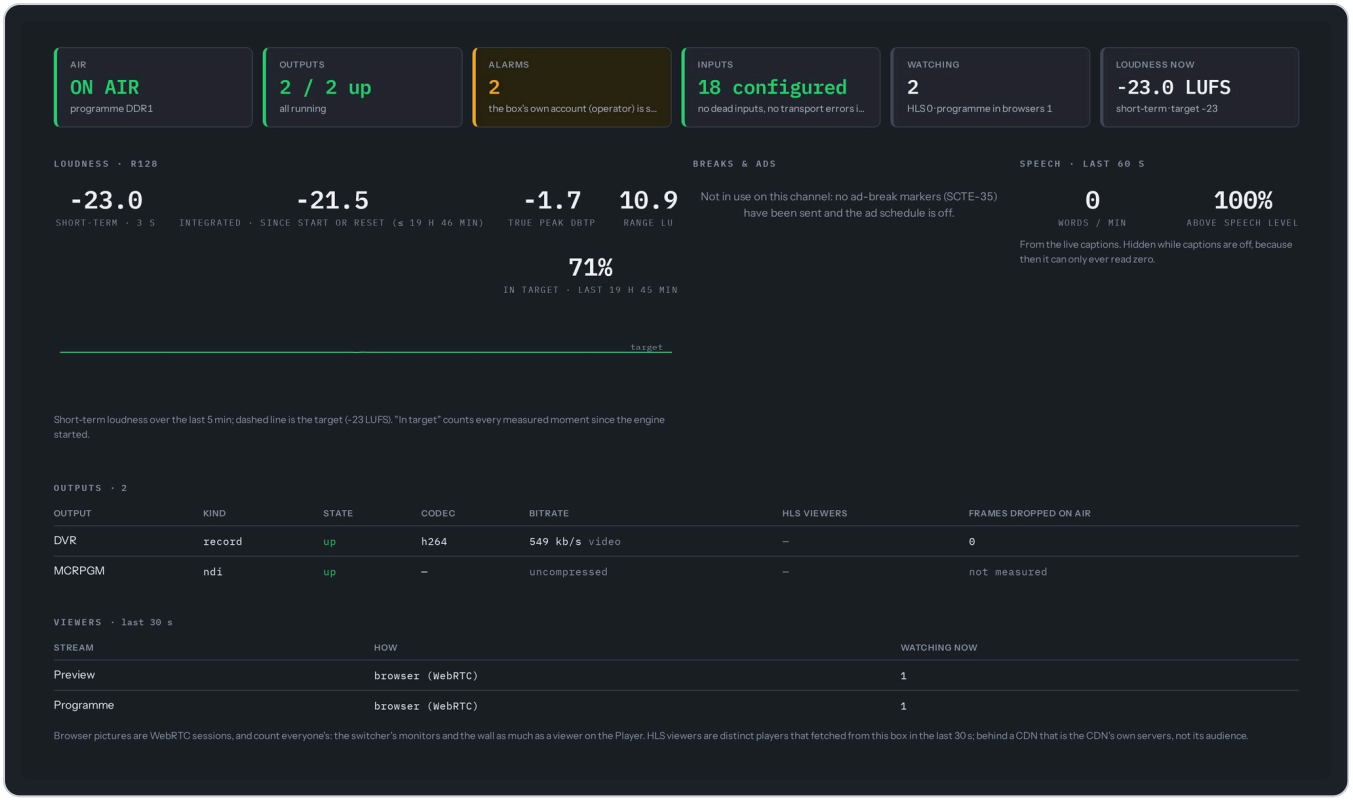
A health strip, every output's bitrate, every input's format and frame drops with TR 101 290 and SRT round-trip and loss, today's as-run with a CSV download, and how many people are watching on the player and HLS.

Loudness compliance

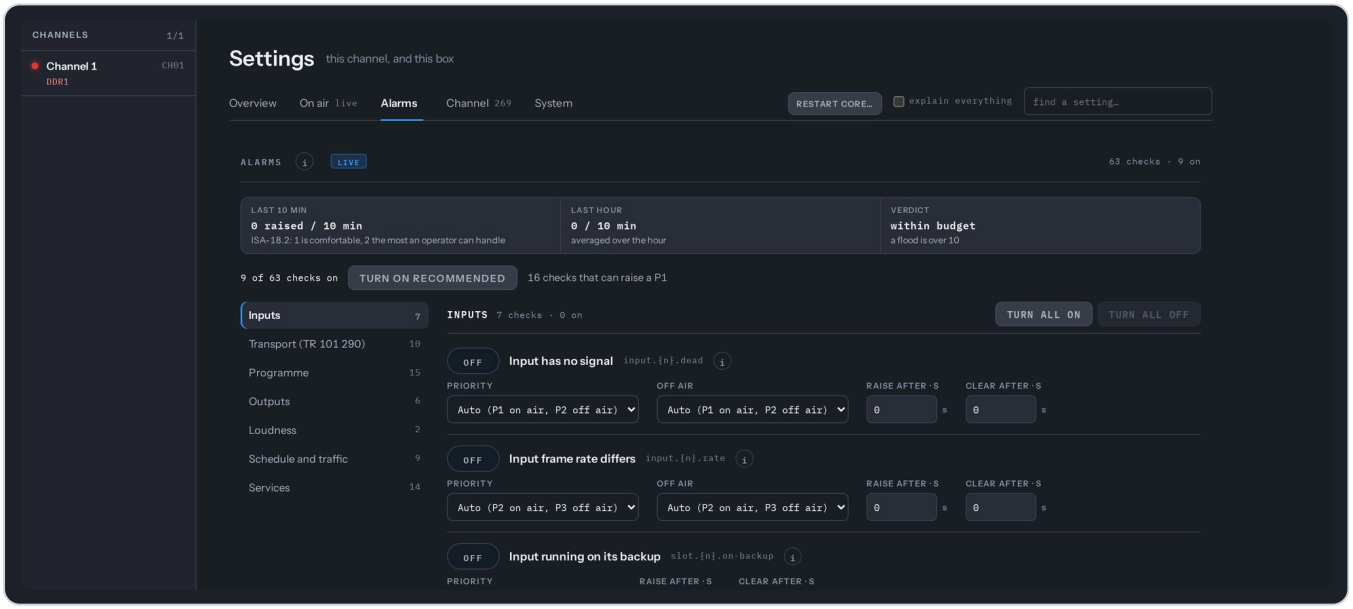
Integrated loudness, range and maximum true peak against the channel's target, with the time in target, recorded as data.

Install and update

One installer stands up any role from a single answers file, with TLS certificates and a firewall. Nodes join with a one-time code and follow the panel's release with signed updates that wait until they are idle.



Analytics. Air, outputs, alarms, inputs, viewers and loudness at a glance, then loudness over time, each output's bitrate and who is watching. demo channel



Settings → Alarms. Every check in groups — inputs, transport, programme, outputs, loudness, schedule, services — each with its priority on and off air and its timing; **Turn on recommended** switches on the checks that can raise a P1. demo channel

20

Security & access control

Secure by default. The panel will not start without a password, and every account belongs to a named person.

Personal logins

A login per person, so the as-run can always say who cut. Repeated failed sign-ins lock the account out; a login can be held without being deleted.

Roles

Admin, engineer, producer, MCR operator, graphics, audio, captions and observer, each reaching only its own pages and commands, and each login or token limited to the channels you tick.

Two-factor authentication

Time-based one-time codes with an authenticator app, required on a new install.

Single sign-on

Sign in through your OpenID Connect identity provider, with verified email addresses only.

Nodes join with a code

Nodes → Add node makes a single-use code and QR, valid for 30 minutes. The new node pins the panel's certificate, so even a self-signed panel on your own network is trusted safely.

Tokens for machines

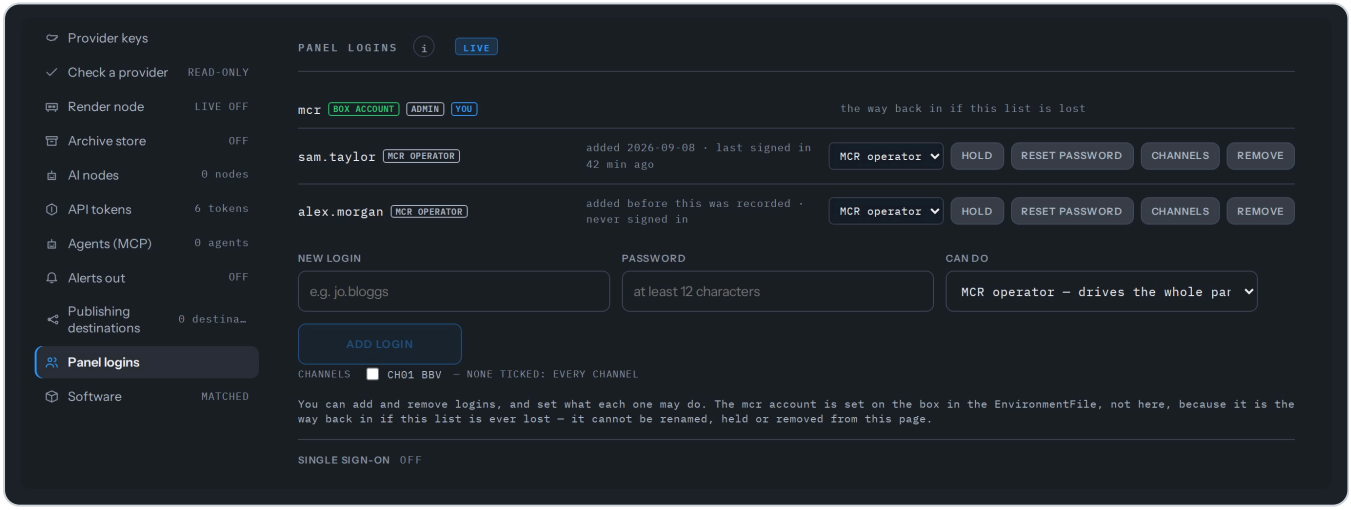
Scoped tokens for render, AI, data and control: each reaches only its own routes, appears in the as-run under its own name, and can be held or revoked.

Secrets stay secret

Stream keys, SRT passphrases and provider keys are write-only in the panel, masked in logs and redacted before storage.

Hardened by default

A new install turns on the firewall, runs services under their own account in a sandbox, sends browser security headers and refuses cross-site writes. Nodes dial out to the panel, so they need no inbound ports beyond their own job.



Settings → Panel logins. A login per person with a role and channel scope; hold, reset or remove; single sign-on below.

demo box · example names

21

Automation, API & AI agents

Drive a channel from hardware panels, automation systems or your own software — and let AI assistants watch it without ever touching air.

REST API

A documented, versioned HTTP API (`/api/v1`) with bearer tokens, behind the same sign-in and roles as the panel — now covering the rundown, PTZ, graphics words, the library's conform, transcribe and translate, settings export and import, and viewers.

Live event stream

Server-sent events push state changes as they happen — including live caption words and graphics state — so a button lights up from the channel itself, not from the press meant to cause it.

Stream Deck and Companion

A Bitfocus Companion module drives Stream Deck, Loupedeck, X-keys and tablet surfaces: cuts, transitions, graphics, replay, breaks, clips, macros and audio description.

Push data in

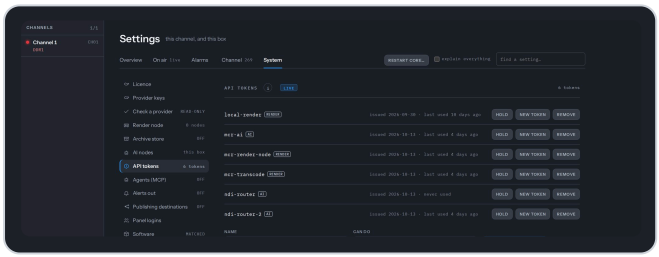
Your scoring, results or newsroom system pushes data to a channel over HTTPS with a data-only token, straight into graphics.

AI agents, observe-only

Connect an AI assistant over the Model Context Protocol (MCP) to read status, alarms, the as-run, outputs and transcript search. Agents can never change what is on air.

Control port

A text control protocol for automation systems, authenticated, bounded, and with every operational command in the as-run.



Settings → API tokens. A token per machine or system, each with what it can do and which channels, when it was last used, and Hold, New token or Remove.

demo channel



Settings → Agents (MCP). "Agents can never change what is on air." An endpoint and a one-line set-up for an agent token.

demo box · example address

22

Scale out: nodes, storage & resilience

Heavy work moves off the air box onto nodes you add as you grow, so air always has the resources it needs.

Render node

GPU or CPU. Renders edits, social versions, house conforms and exact clips, checks every output, and keeps captions through the render.

Transcode node

Makes live SRT, RTMP(S), HLS and NDI outputs from one mezzanine stream sent by the channel, on the network where they are needed.

AI node

Speech, translation, voices, speakers, picture tags and summaries for several channels from one machine with an NVIDIA graphics card.

Ingest node

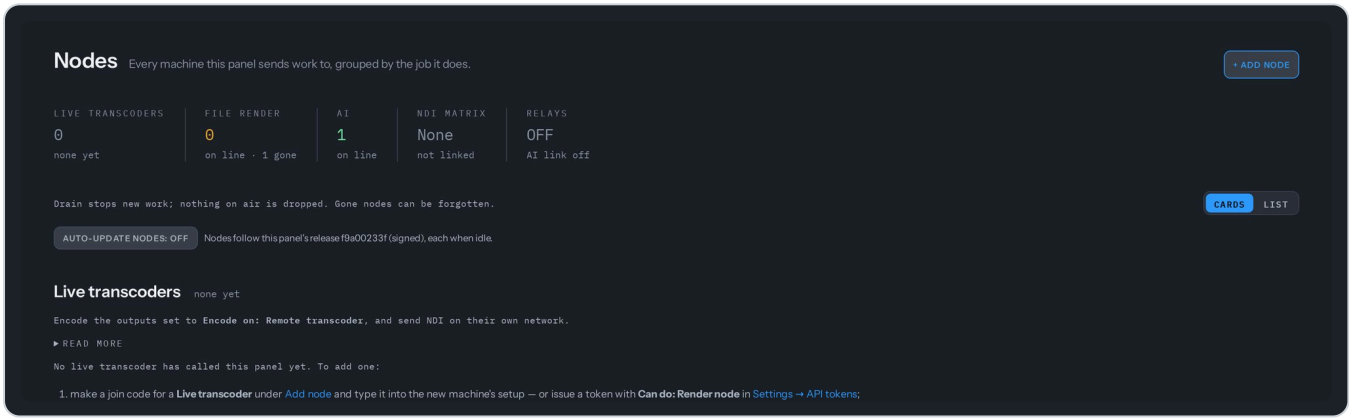
Turns contribution feeds into NDI and relays one SRT feed to several channels.

NDI matrix BETA

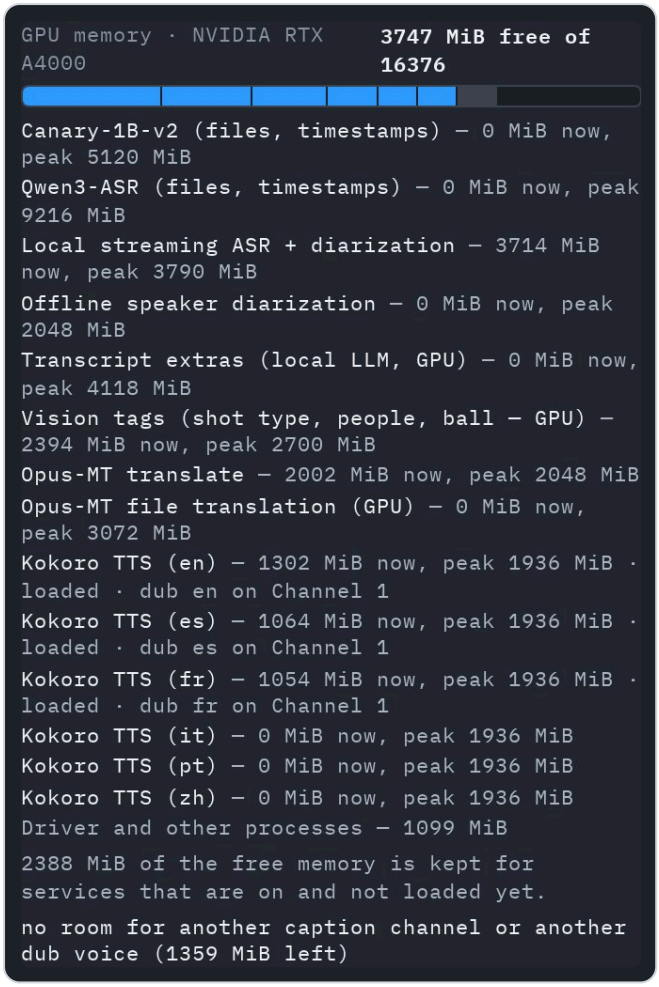
An NDI router node whose own control page opens inside the panel as **Matrix**, with its outputs in every channel's NDI source list.

Shared media library

One library for channels and nodes on NFS, SMB or a SAN, with S3 transfer by short-lived links.



Nodes. Every machine the panel sends work to, grouped by job, with its health, version and current work; automatic updates for all of them, or one at a time. demo box



Nodes → GPU memory. Each service on the card, what it holds now and its peak, then the room left in words. demo box

Two-core resilience

Two channel cores run the same channel in lock-step on the wall clock, every command mirrored to both, with an edge that cuts between them — or a warm standby that takes over. It is in the software and being prepared for first live service; ask us before planning around it.

Drain and grow

Take a node out of service without dropping anything on air — a drain that would take something off air is refused — and start with one box running everything, adding nodes as channels and services grow.

23 Formats & standards

AREA	SUPPORTED
Contribution in	SRT (caller / listener, AES), RIST, RTMP / RTMPS / Enhanced RTMP, NDI (with Discovery Server and Access Manager), HLS, WHEP / WebRTC, browser guests
Video in	H.264, HEVC, MPEG-2; AV1 where the decoder is installed; decoded on the CPU, NVIDIA or Intel
Camera control	VISCA over IP and VISCA over TCP (beta)
Transport checks	TR 101 290 priority 1 and 2 on SRT and RIST inputs
Delivery out	SRT (MPEG-TS, strict CBR option), RTMP, RTMPS, Enhanced RTMP (HEVC, AV1), HLS with ABR ladders, NDI (local or from a transcode node), WebRTC (WHEP) monitoring, file recording; programme, clean feed or multiview
Video out	H.264 and HEVC; AV1 on Enhanced RTMP and file renders; MPEG-2 and 1080i legs; 720p or 1080p programme
Encoders	Software (CPU), Intel Quick Sync, NVIDIA NVENC, VA-API; chosen per channel and per output
Audio	AAC; AC-3 (DVB or ATSC signalling); E-AC-3 and SMPTE 302M; stereo or 5.1 programme with ITU-R BS.775 fold-down; multi-track MPEG-TS with fixed PIDs and ISO 639 labels; multichannel NDI as pairs or strips
Frame rates	23.976 – 60 fps including 29.97 and 59.94; 100 / 120 fps high-frame-rate replay
Captions	CEA-608 / 708 in H.264 and HEVC (ATSC A/53) with PMT signalling, in outputs or in the air encode, authored or passed through; T.35 in AV1 file renders; WebVTT; HLS subtitle renditions; SRT and VTT files; burnt-in, including translations and bilingual pairs
Languages	28-language catalogue for live recognition. Local translation from English: 7 languages. Local voices: 8. Others through cloud services. Languages CEA-608 cannot carry go to subtitles and burn-in
Loudness	EBU R128 (-23 LUFS), ATSC A/85 (-24 LKFS); momentary, short-term, integrated, LRA, true peak
Ad signalling	SCTE-35 splice insert and time signal with segmentation; HLS cue-out / cue-in and date-range tags; upstream SCTE-35 relayed
Programme guide	CSV, XMLTV and BXF import; XMLTV and Roku feeds; Samsung TV Plus and Pluto TV exports; DVB EIT present / following
Render profiles	9:16, 1:1, house master (H.264, HEVC, AV1), house-format conform, DNxHR and AVC-Intra 100 masters
Live data	Google Sheets, CSV, JSON and RSS by address; push over HTTPS
Publishing	S3-compatible storage; YouTube, TikTok, X, LinkedIn, Facebook, Instagram, Threads (beta)
Control	Web panel, REST API with server-sent events, Bitfocus Companion, macros, text control port, MCP (observe-only)
Access control	Personal logins with lockout, eight roles, TOTP two-factor, OpenID Connect SSO, scoped API tokens, node join codes, audit trail

24 How it behaves, and what is next

Designed to tell you, not to guess

- It never cuts away from a bad source by itself: armed failover and the optional slate backstop are the only source changes it makes to air on its own.
- It never decides on its own that a silent feed is wrong. It raises an alarm and waits for you.
- It never opens a firewall port to the internet for you beyond the ones its features use.
- Captions, voices, comms, recording, uploads and a failed output are all contained. None of them can take the programme off air.
- The one exception is the licence: at the end of a trial, or of a lapsed licence's grace period, outputs show a notice until a key is entered.
- AI agents can watch; only people change what is on air. Replay ready can cue an idle replay channel, but only a person plays it.

On the roadmap PLANNED

- Two-core resilience in live service
- Ducking a follow mix under its voices automatically
- PTZ through a remote site, ONVIF and NDI PTZ, and tally to the camera
- Translating file transcripts from languages other than English
- Delivery presets that enforce, not only warn
- Social publishing proven against each live platform

Roadmap items are plans, not commitments, and may change.

Upgrading from release 0.1.0? Defaults that changed

- Loudness correction and the loudness alarms are on.
- The slate backstop is off until armed in Air safety.
- Fade to black takes keyed graphics down too.
- Browser monitors follow the channel's size and rate, which costs more CPU above 1280 wide or 25 fps.
- 5.1 folds to stereo at BS.775 levels, about 10 dB louder: recheck 5.1 trims.
- A multichannel feed with no layout airs channels 1+2.
- Archive recordings keep 90 days, ISO archives 7.
- Off-shape files are letterboxed, not stretched.
- Dubs are placed like an interpreter: late lines are spoken late.
- Translations arrive a sentence at a time.
- New DVE channels build all eight boxes.
- Local voices and summaries need an NVIDIA GPU.
- Graphics is two pages, Make and Play; Design and the Workspace are gone.
- New channels dissolve for one second; an upgraded channel keeps its rate.
- A layout or scene taken with a transition now dissolves a box that changes source.
- A scene with a main picture, clips or keys goes up only through its own take, not from a schedule or macro.

WHAT IS NEW SINCE RELEASE 0.1.0

Sound & access
dubs on the desk · interpreter placement · follow mixes · international clean · language detection · burnt-in translation · 28-language catalogue · word-by-word player captions · voices that load on demand · local captions per channel

Production
Make and Play · web pages on the Board · scene ARM · box dissolve and swap · 23 animated templates · NDI graphics · PTZ · M/E bank · frame-accurate replay angles · Highlights page · shot search · Replay ready

Playout & delivery
rundown builder · house-format QC and conform · captions in the air encode · multiview as an output · NDI from a transcode node · encode and decode choice

Operations
AI GPU budget · hardware verdict · analytics and viewers · settings export and presets · restart preview · node join codes and signed updates · NDI matrix · 30-day trial

Credits & notices

Example pictures

Screens are from a demo channel. Every video picture shown in them is an example frame from Wikimedia Commons, cropped and resized, placed in the browser for these pages only:

- “Dawn of Darkness — Live Closing Passage”, Mortaz Band, by Mostafa Meraji — CC0 (public domain dedication).
- “Aerial imagery of Zug Switzerland” by anotherview — CC BY 3.0, creativecommons.org/licenses/by/3.0.
- “Close-up fireworks during Lunar New Year celebration in Zhongshan, China, 2026” by JaydenChao — CC BY 4.0, creativecommons.org/licenses/by/4.0.
- “Cars drifting around a race track at twilight and night as viewed from FPV drone” by anotherview — CC BY 3.0, creativecommons.org/licenses/by/3.0.
- “Rome marathon, via Galvani, 2017” by Husky — CC BY 4.0, creativecommons.org/licenses/by/4.0.
- “Colorful Autumn — Stage Excerpt”, Mortaz Band, by Mostafa Meraji — CC0 (public domain dedication).

The rundown screen lists public-domain films. The live-caption pictures show NASA footage (public domain; NASA does not endorse this product), played through a demo channel with captions switched on; their caption text is what the recogniser wrote live.

Other caption and transcript text, guest names and graphics text in the screens are examples. Login names, machine names and network addresses have been replaced.

Trademarks

AC-3 and E-AC-3 are audio coding formats named here by their standard designations. NDI, SRT, NVIDIA, Intel, YouTube, Twitch, TikTok, X, LinkedIn, Facebook, Instagram, Threads, Samsung TV Plus, Pluto TV, Roku, Amazon, Wurl, Stream Deck, Companion, Contour, ShuttlePRO, VISCA and every other product or company name are trademarks of their owners, used only to identify compatible products and services. MCR is not affiliated with or endorsed by any of them.

About this guide

Describes MCR as of November 2026, the software after release 0.10. Features and their limits are as described; items marked beta, optional or NVIDIA GPU are described as they are. Nothing in this guide is a contractual specification; the licence agreement sets the terms of use.

M C R

Live production, master control, captions and delivery — in a browser.

SEBIU LABS
CONTACT@SEBIULABS.CO.UK
EDITION 4.1 · NOVEMBER 2026