STATE OF THE UNION

GStreamer Conference

07 October 2024, Montréal

Tim-Philipp Müller < tim@centricular.com >

gstreamer

Centricular

WELCOME

Nice to see everyone again in person!

BEFORE WE GET STARTED ...

JUST THE OTHER DAY

Wim Taymans presented "20 Years of GStreamer" at the GStreamer Conference in Lyon in 2019.

That was 5 years ago!

HAPPY 25TH BIRTHDAY GSTREAMER!

image by lamarogre

TESTAMENT TO
OUR AMAZING COMMUNITY!

Speaking of which ..

MAILING LIST TO DISCOURSE

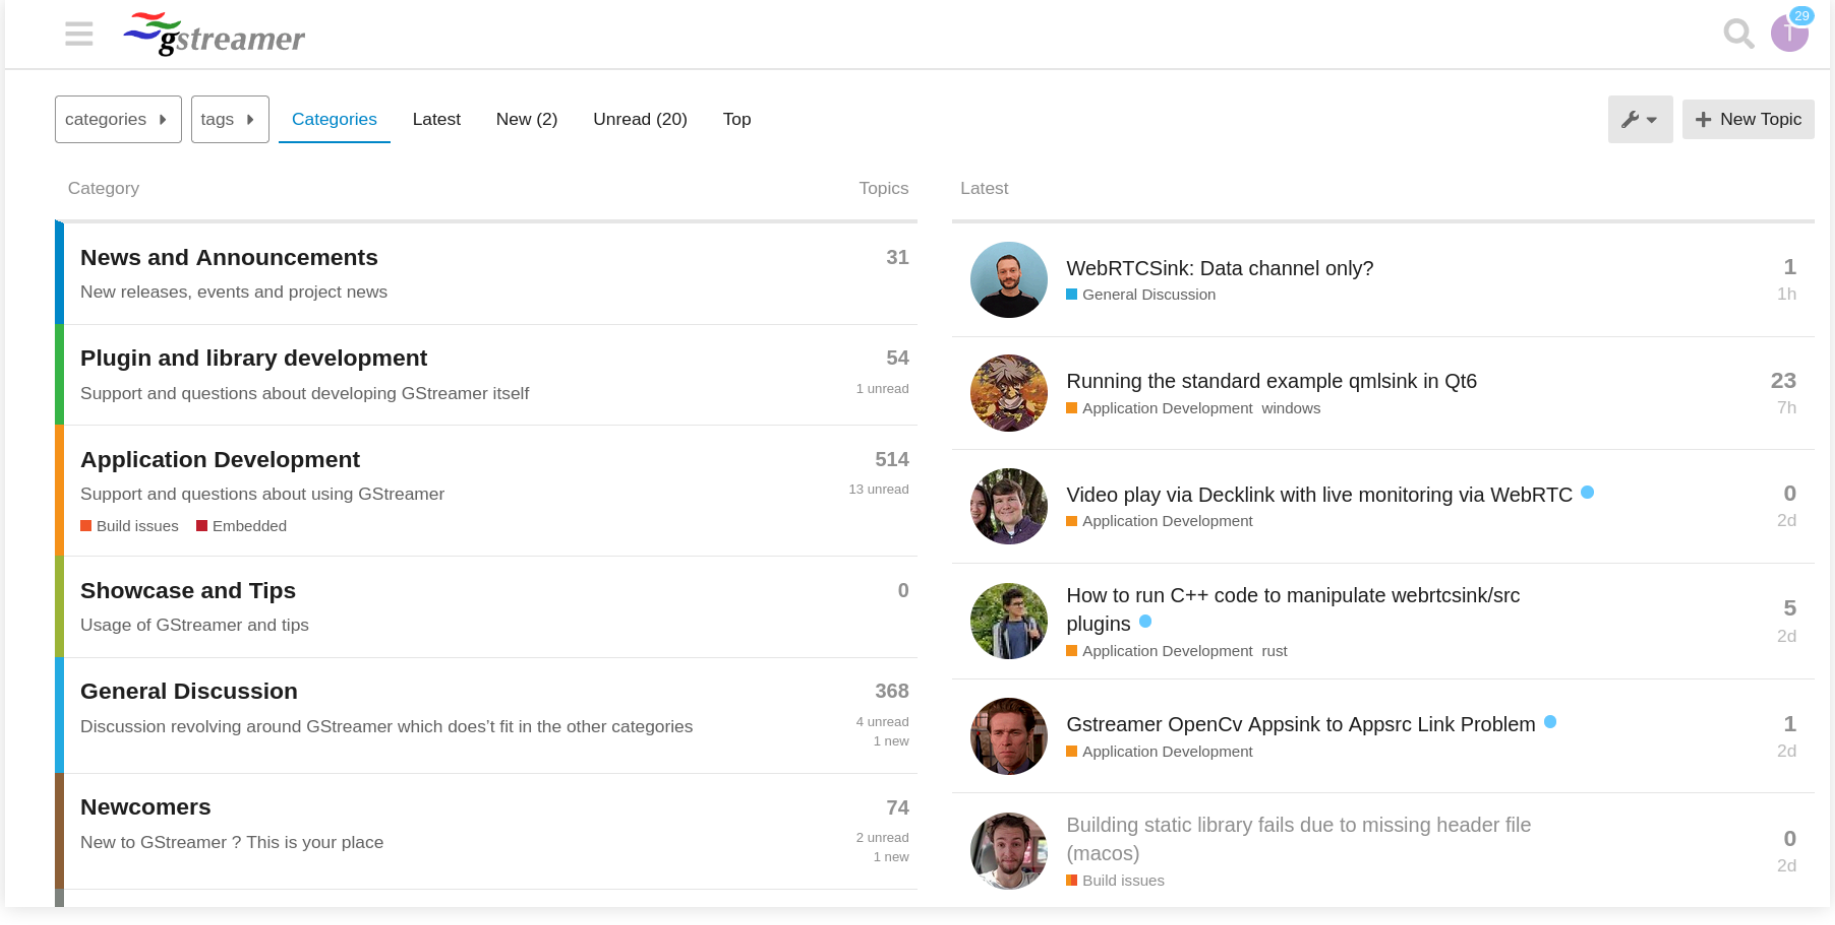

DISCOURSE STATS

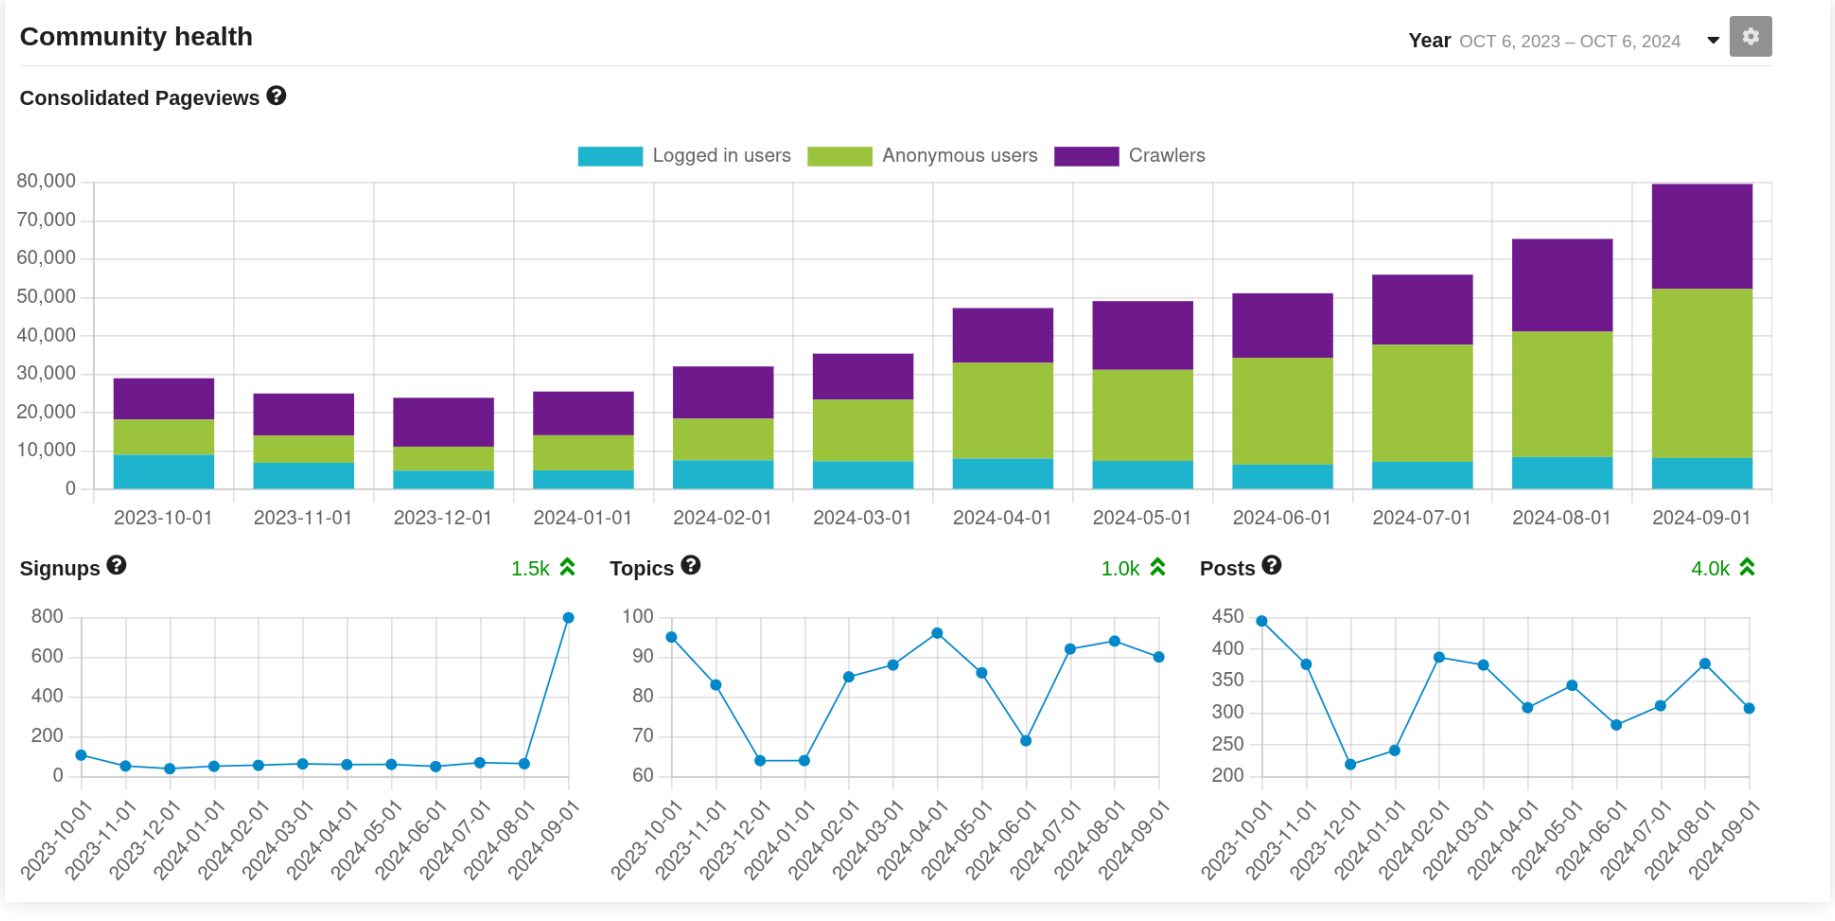

IRC TO MATRIX

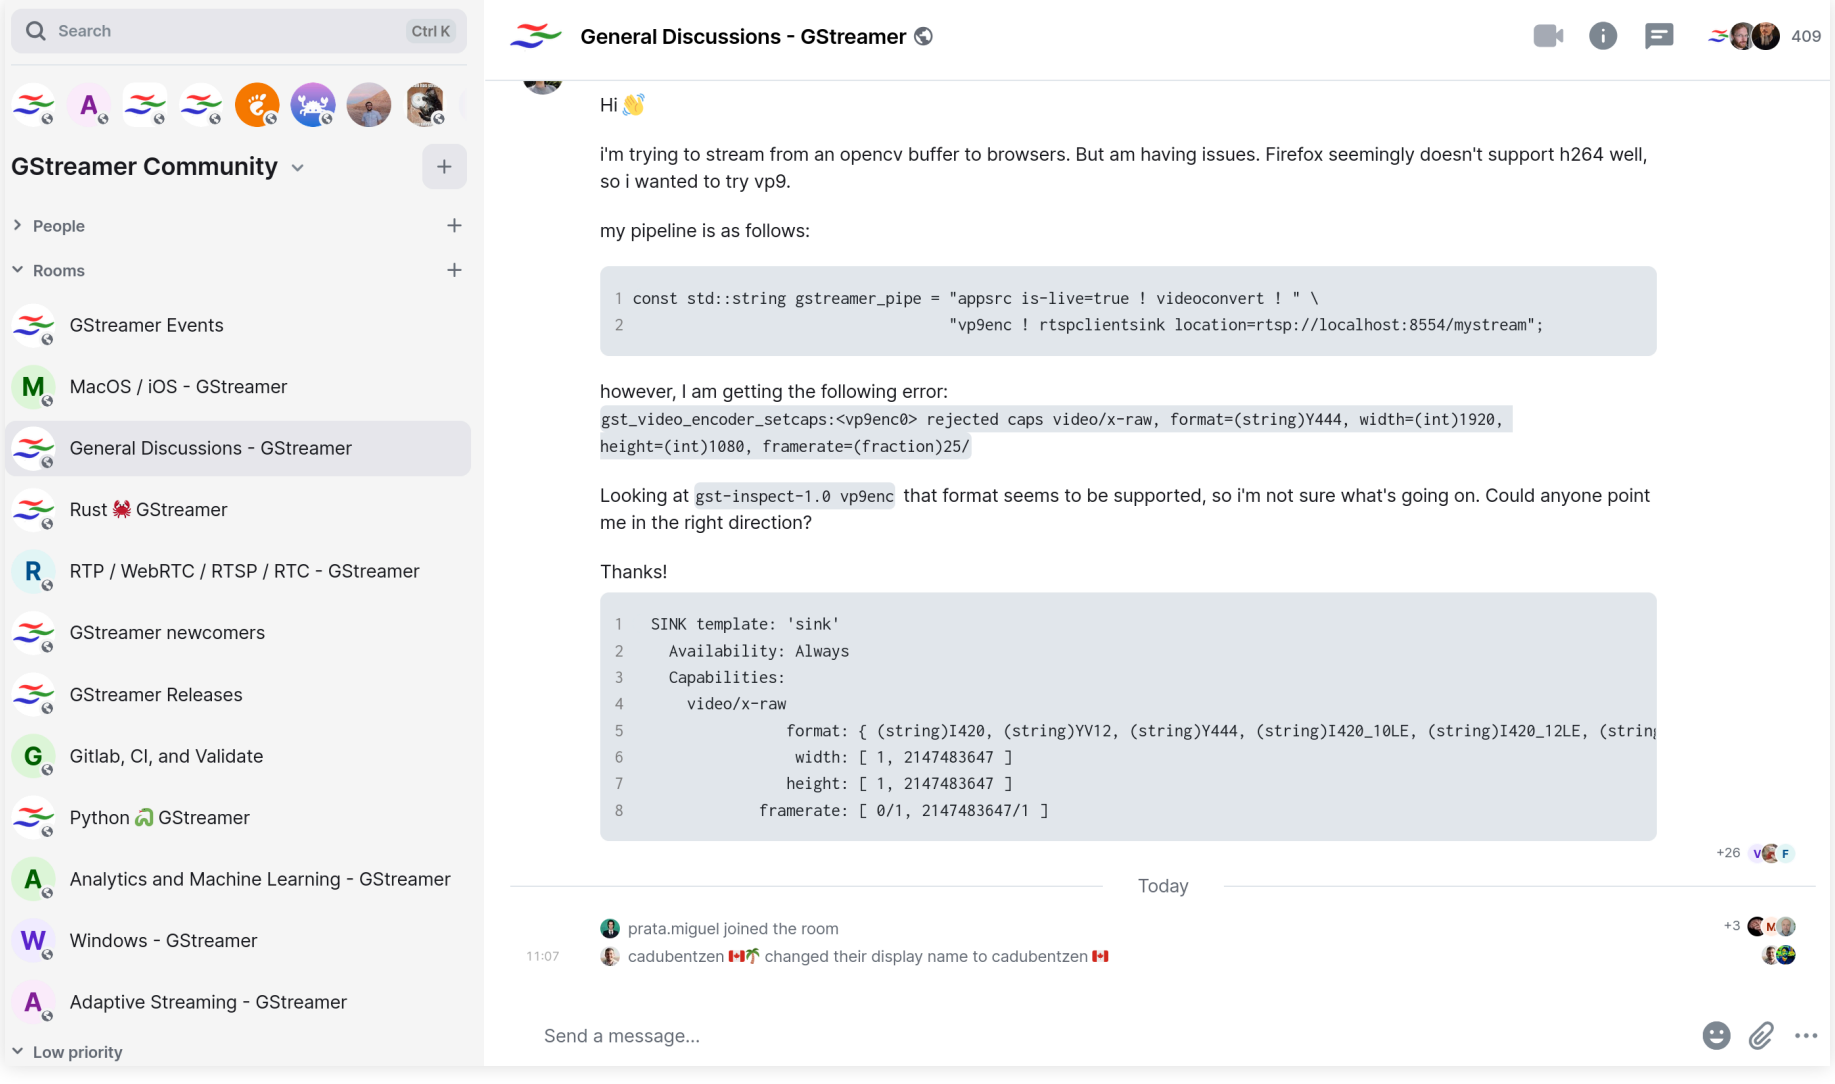

WHAT ELSE HAVE WE BEEN UP TO?

RELEASES!

- 1.20: February 2022
- 1.22: January 2023
- 1.24: March 2024
- 1.26: ???

1.24 STABLE SERIES

- 8 bug-fix releases since March
  - 800 commits, 450 backported merge requests
  - Every ~3 weeks

UPCOMING 1.26

- ~2000 commits, 1400 merge requests (so far)
  - ca. 15-20% of commits, 35-45% of MRs are Rust!

LET'S TALK A LITTLE ABOUT RUST!

"Fast, safe and productive - pick three."

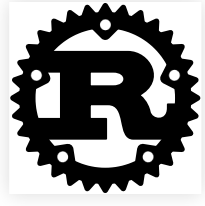

WHY RUST?

- Perfect language for us technically.
- Excellent C compatibility.
- Fantastic community.
- Superb ecosystem.
- Adoption by all the major industry players.

THE FUTURE OF GSTREAMER IS RUST

- **Bindings**: Mature and battle-tested
- **Plugins**: Where cool™ new stuff happens
  - Can be used from any language
  - Shipped for macOS, Windows, Android, iOS
  - Linux distros still catching up

THE FUTURE OF GSTREAMER IS RUST

- Don't worry!
  - Change will be gradual, C isn't going anywhere.
- However!
  - Rust plugins are needed for a fully-featured GStreamer

BACK TO OUR FEATURE PRESENTATION..

RTP + RTSP STACK REWRITE IN RUST

Sponsored by the Sovereign Tech Fund

Why? Security and Long-Term Maintainability

RTP SESSION HANDLING IN RUST

- rtpsend + rtprecv = rtpbin2
- Low-level RTP/RTCP packet handling all in Rust too
- All the basic features are there
- Advanced feature still todo

NEW RTSP SOURCE IN RUST

- rtspsrc2
- Live streaming only for now
- Can use new or existing RTP stack

NEW RTP DEPAYLOADERS/PAYLOADERS IN RUST

- New base classes in Rust
- Packet parsing/handling all in Rust
- Please give them a spin!

MISC RTP IMPROVEMENTS

- RFC7273 clock sync support (PTP, NTP sync)
- Easy way to get original timestamps from rtpjitterbuffer
- `rtppassthroughpay` element
  - useful for RTP passthrough in rtsp-server
  - regenerate RTP timestamps based on buffer ts

RTSP CLIENT SOURCE

- Optionally timestamp RTP packets with their receive times in TCP/HTTP mode
- Spec compliant control url handling
  - Wrongly implemented in many places
  - Automatic fallback for non-compliant servers
  - Property to force non-compliant urls if needed

WHAT ELSE? WEBRTC!

WEBRTC: WHAT IS IT ABOUT?

- *"How do I stream to a web browser?"*
- Low latency, works pretty much everywhere.
- Different codecs
- Leverage all of GStreamer:
  transmuxing, hw-acceleration etc.

WEBRTCBIN (LOW LEVEL)

- Better negotiation
- Better spec compliance

WEBRTC MADE EASY

WEBRTCSINK + WEBRTCSRC

- Batteries-included WebRTC sender + receiver:
  - codec negotiation
  - congestion control
  - automatic bitrate adaptation

WEBRTC MADE EASY - CONT'D

WEBRTCSINK + WEBRTCSRC

- Signallers and Associated Elements
  - LiveKit, Amazon Kinesis, Janus VideoRoom
  - WHEP, WHIP
- New webrtcsrc to complement webrtcsink

WEBRTC MADE EASY

NEW WEBRTCSINK + WEBRTCSRC FEATURES

- Support for already-encoded streams
- Raw payload support (raw audio, video)
- More encoders
- Built-in embedded signalling / web server

MOVING ON TO OTHER TOPICS..

WEB TECHNOLOGIES

- Media Source Extension (MSE)
- Encrypted Media Extensions (EME)
- Wasm

ANALYTICS + MACHINE LEARNING

- New GstAnalyticsRelationMeta
  - observations and relationships
- Object detection, classification and tracking

ANALYTICS - CONT'D

- ONNX elements
  - onnxinference
  - ssdobjectdetector (to be split out)

ANALYTICS - CONT'D

- Analytics overlay
  - visualises objects
- originalbuffersave, originalbufferrestore
  - Restore the original buffer after analysis on a transformed version

PLAYBACK

- playbin3, uridecodebin3, parsebin
  - better stream selection, gapless handling, buffering
  - many corner case stability/reliability issues fixed
  - Todo: allow decoder selection tweaking
- GstPlay library + gst-play-1.0
  - use playbin3 by default now

CAPTURE + PLAYOUT CARDS

- **AJA** sink + source upstream now
  - incl. device provider
  - HANC/VANC ancillary data support
- **Decklink**
  - HDR output and input (PQ + HLG)
  - more modes, devices
  - better frame scheduling

PIPELINE SPLITTING

- For robustness + easier handling
- Multiple pipelines, either:
  - within the same application
  - in different processes on same machine (IPC)
- Typically 1:1 or 1:N

PIPELINE SPLITTING: INTRA-PROCESS

- New inter plugin for 1:N and any kind of data
  - inspired by gst-interpipes
  - some restrictions for now

PIPELINE SPLITTING: INTER-PROCESS

- **Send video or data to another process (IPC):**
  - New **unixfdsink**, **unixfdsrc**
    - uses memfd + dmabuf
  - Windows **D3D11** + **D3D12** video
  - **CUDA** video sink + source
- **GstMeta serialisation + deserialisation**
  - for certain common metas
  - used in unixfd and CUDA elements

CLOSED CAPTIONS

- cea608tocea708
- tttocea708
- cea708mux and cea608mux
- cea708overlay
- cea708 insertion in more encoders
- caption generation improvements in transcriberbin

SPEECH TO TEXT: TRANSCRIPTION + TRANSLATION

- Translation support in transcriberbin
  - awstranscriber gained translation support
    - multiple languages at the same time
- New Speechmatics transcriber
  - includes translation support
  - can be deployed on-premises, unlike aws

ANCILLARY DATA

- GstMeta for SMPTE ST-291M HANC/VANC Ancillary Data
- SMPTE ST-2038 ANC support in MPEG-TS
- ID3 in MPEG-TS

TRACING + DEBUGGING

NEW TRACERS

- **Pad push durations** tracer
  - measures the time it takes for a buffer/buffer list push to return
- **Save pad data flow** into a .pcap file

NEW CODEC SUPPORT

- **LCEVC: Low Complexity Enh'ment Video Coding**
  - enhancement on top of different base codec layer
  - decoder and encoder
- **JPEG XS**
  - low latency, for production workflows
  - MPEG-TS container support implemented
  - Todo: RTP, MP4, MXF

VIDEO FORMATS

- More video pixel formats
- More software conversion fast paths
- More formats supported in OpenGL, D3D11, D3D12
- Bayer 10/12/14/16-bit depth support

HARDWARE-ACCELERATED VIDEO

- **DRM modifiers support** for Linux dmabufs
  - plus negotiation, handling in va, wayland, msdk
- waylandsink **DRM Dumb allocator support** + 10bit support
- Lots of work on Vulkan and Vulkan video codecs support

VIDEO4LINUX

- Lots of improvements all over the place
- Stateful and stateless codec support
- New stateless AV1 decoder
- Stateless decoders tested on CI using QEMU + visl
- Encoder support for keyframe requests
- Stateful decoders HDR10 support

VA-API SUPPORT

- va replaces gstreamer-vaapi
  - Based on shared codec base classes
    - Just works™ \o/
    - Decoders have PRIMARY+1 rank
- New AV1, VP9, JPEG encoders
- More rate control modes (ICQ, QVBR)
- Performance improvements
- Support for more pixel formats

NVIDIA + CUDA

- Desktop GPU decoders + encoders rename
  - `nvh{h264,h265,vp8,vp9}sldec` to `nv{h264,h265,vp8,vp9}dec`
  - `nvcuda{h264,h265}enc` to `nv{h264,h265}enc`
- AV1 encoder added
- Encoder RGB formats support
- CUDA stream integration support
- D3D12 integration
- nvCOMP plugin for lossless compression on GPU

AUDIO

- ASIO: build without SDK headers
  - always ship the plugin
- DSD audio support
- LC3 bluetooth codec

MACOS

- audio sink source improvements, latency etc.
- gst-validate support (gst_macos_main, lldb)
- vtenc hardware encoder improvements
  - incl. HEVC alpha encoding support
- Apple AAC encoder (atenc)

ANDROID

- Media CODEC ported to the native API
  - reduces amount of Java <-> native calls = better performance
- AV1 decoder/encoder
- Rust plugins shipped in binary packages

WINDOWS

- New Microsoft WebView2 based web browser source
- DirectWrite text rendering
- D3D12 plugin, support and integrations
- d3d12swapchainsink for Windows composition API based apps
- PTP clock support is now also available on Windows
- Many many other improvements

UI TOOLKIT INTEGRATION

GTK4 PAINTABLE SINK

- performance + integration improvements
- dmabuf import support
- support for rotations / flipping / scaling
- fullscreen property
- black background setting
- GL support on Windows

UI TOOLKIT INTEGRATION

QT6 ELEMENTS

- OpenGL src, sink, mixer, overlay
- D3D11 sink for Windows

EDITING SERVICES LIBRARY

- Reverse playback support
- encodebin2 support
  - can use a muxing sink, e.g. HLS/DASH/splitmux

STREAMING: HLS

- hlscmafsink for serving fragmented MP4 in HLS
- Lots of Rust fmp4 muxer improvements
  - incl. support for AV1, VP8, Opus, FLAC
  - better handling of caps changes

STREAMING

- New **QUIC source and sink** (quinnquicsrc, quinnquicsink)
- New **mpegtslive element** that wraps an existing live MPEG-TS source (udp, srt)
  - Provides a clock based on the PCR of the stream

MISCELLANEA

- FFmpeg 7.0 support
- ORC AVX2 support
- splitmuxsrc/sink: dynamic fragment addition
- jpegparse has a rank and is now autoplugged
- wpesrc: WPEWebkit 2.0 support

THE FUTURE ...

MORE RUST REWRITES?

- Demuxers
- Parsers
- WebRTC? (ICE wip)
- RTSP server?

ANALYTICS, ANALYTICS, ANALYTICS!

MACHINE LEARNING!

Stay tuned

RE-ORGANISE MODULES?

Update and simplify our module story

CONTINUOUS INTEGRATION

- Run tests and integration testsuite
  on Windows + macOS
  - validate gained macOS support
- Test hardware acceleration support in upstream CI
- More sanitisers (ASAN, USBSAN etc)

NIGHTLY BUILDS

- Has been on the list for a while
- More infrastructure work required
- Goal: create release binaries on the CI
- Signing of binaries would be nice too

TRACING + DEBUGGING IMPROVEMENTS

- more tracers
- nice UIs
- improve output side (common log/trace formats)

IMPROVE PIPEWIRE INTEGRATION

Existing plugins not in great shape

HLS/DASH AUTHORING IMPROVEMENTS

- Master playlist management for HLS alternate renditions
- Low-latency HLS

THAT'S ALL!

THANK YOU!

QUESTIONS?