

Tvheadend - Feature #6061

Expose and pass information to commands in "MPEG-TS Spawn" from TvH

2021-06-04 01:08 - Bogdan I.

Status:	New	Start date:	2021-06-04
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:	Streaming	Estimated time:	0.00 hour
Target version:			

Description

Commands run by "MPEG-TS Spawn" Streaming Profiles could benefit from injecting basic data from Tvheadend.

For example, one could export the profile name, service-id, channel-id etc. by using `setenv()`, or by doing string replacement.

Setting Environment Variables via `setenv()` would probably be simpler to implement than string replacement and it's probably safer for the end-user.

For example, doing this in `src/spawn.c`:

```
setenv("THISISNEW", "test", 1);
```

Will properly set an environment variable called `THISISNEW="test"`.

When creating a streaming profile that uses a bash script with the following contents:

```
#!/bin/bash
export -p > "/tmp/dump-vars"
/usr/bin/ffmpeg -nostats -v verbose -i - -c copy -copy_unknown -f mpegts -
```

You can find the `THISISNEW="test"` in `tmp/dump-vars`

Unfortunately my C skills are pretty much non-existent, and my attempts in actually grabbing the info and passing them to `spawn_with_passthrough()` have failed spectacularly.

What would be the use-case?

As for the motivation for this feature, I'm currently limited to ~2GiB GPU memory on my Tvh machine that runs with a cheap GTX 1050. Due to the current GPU shortage we are facing right now, upgrading to anything with more VRAM doesn't look too possible in the near future...

Loading up a full 1080p/i stream in GPU memory to transcode uses about 400-450MiB VRAM.

Unfortunately, if 2 users are currently watching the same channel/stream, even with the same Streaming Profile, the transcoding operation is done twice, thus using resources twice, needlessly.

In my quest to write a script to tackle this, which would write HLS and then re-read it if the Stream is already transcoding... I realized that I'm kind of limited by not having an actual way of know if a specific stream is already being transcoded.

Ideally I should be able to create a `stream_profile_name-stream_id.pid` file from my script to check if there's a transcode going on for that specific stream with that specific profile

This could probably be useful for lots of people who are adventurous enough to implement their own transcoding scripts/tools.

History

#1 - 2021-09-25 15:17 - Patrick Brauer

+1 having some basic information about the incoming stream data would also help a lot to reduce latency if these allow us to skip analyzing the stream (again) on our own once passed to us. just by letting us know which service is subscribed we could lookup most of that using the API but we don't know anything about the stream except for whats described in the stream itself.