The Portal has two audio problems. The first is that it can't always be heard. For communication that's primarily aural, this is a serious problem. The second issue is that the production quality is inexplicably inconsistent and this compounds the first problem.

The quality of a podcast's audio shouldn't draw attention to itself. Ideally, the listener should set their device's volume at the start of an episode and not touch it again. Achieving this requires success at five levels of analysis.

1. well-recorded audio,
2. well-processed audio,
3. consistency within an episode,
4. consistency between episodes, and
5. consistency between different podcasts.

Well-recorded audio. Recording good audio depends on the source (people speaking), environment, selecting the right tools for the job and employing them appropriately. If audio is recorded poorly, it's not always possible to save it in post-production. When it is, it's often time consuming.

Well-processed audio. Exactly what this goal entails and how long that takes will depend on the quality of the recording. Generally you want to get rid of unwanted noise and ensure a relatively consistent perceived loudness. Perceived loudness is measured in Loudness Units (LU).

Like the recording process, the end product of audio processing is affected by the engineer's source material (recorded audio), listening environment, the tools they use and their capacity to use them. It's hard to fix a problem you can't hear.

Consistency within an episode. Most podcasts consist of several segments: intro, outro, conversation and advertisements. The audio's quality and perceived loudness should be relatively consistent from one segment to the next.

Consistency between episodes. Assuming there's consistency within an episode, that consistency should be maintained between episodes.

Consistency between different podcasts. Attempts have been made to achieve this via recognised standards of average perceived loudness, aka integrated loudness or LUFS (Loudness Units relative to Full Scale).

Unfortunately, different platforms apply different standards. Spotify is -14 LUFS, YouTube is -13 LUFS, the list goes on. Unless you want to master each episode for every platform's standard, it's best to use the Audio Engineering Society (AES) standards for podcasting — no more than -16 LUFS and no less than -20 LUFS. From that standard we can infer an average of -18 LUFS.

So, where does The Portal fall short? At every level except recorded audio. Don't get me wrong, there's room for improvement with regard to recorded audio, but it's not bad enough to care..for now.

Audio Processing

When explaining well-processed audio, I mentioned getting rid of unwanted noise. That happens during The Portal's episodes, but only sometimes.

Pronouncing "T" and "P" into a microphone causes it to be hit with a gust of air which translates to bass and sub-bass frequencies in the recorded audio. The former can be heard, the latter can only be felt. For most laptop speakers and cheaper earphones/headphones, this isn't an issue because they can't reproduce

those sounds. However, if you're listening in your car, home stereo, or computer monitors (depending on type), you're more likely to hear those frequencies - and they're not subtle.

Those bass and sub-bass frequencies can be minimised or removed entirely during processing. Unfortunately, this is only occurring in some segments of The Portal, in some episodes, some of the time.

The other noise you want to deal with is sibilance, caused by pronouncing "S," "SH" and "Z." They result in harsh, higher frequencies akin to scraping fingernails on a chalkboard. These can't be removed entirely without negatively affecting the end product, but they can be minimised. Again, this is only occurring in some segments, some episodes, some of the time.

The second goal for audio processing is relatively consistent perceived loudness within a segment. That relative consistency is known as Loudness Range (LRA), measured in LU (Loudness Units). An LRA of 8 LU or less is a good ballpark figure.

The graph below shows the LRA of individual segments from the last ten episodes. Unfortunately, initial data points from episodes 17, 22 and 23 are such outliers that they've fucked up the scale of my graph, making it difficult to distinguish other data points.

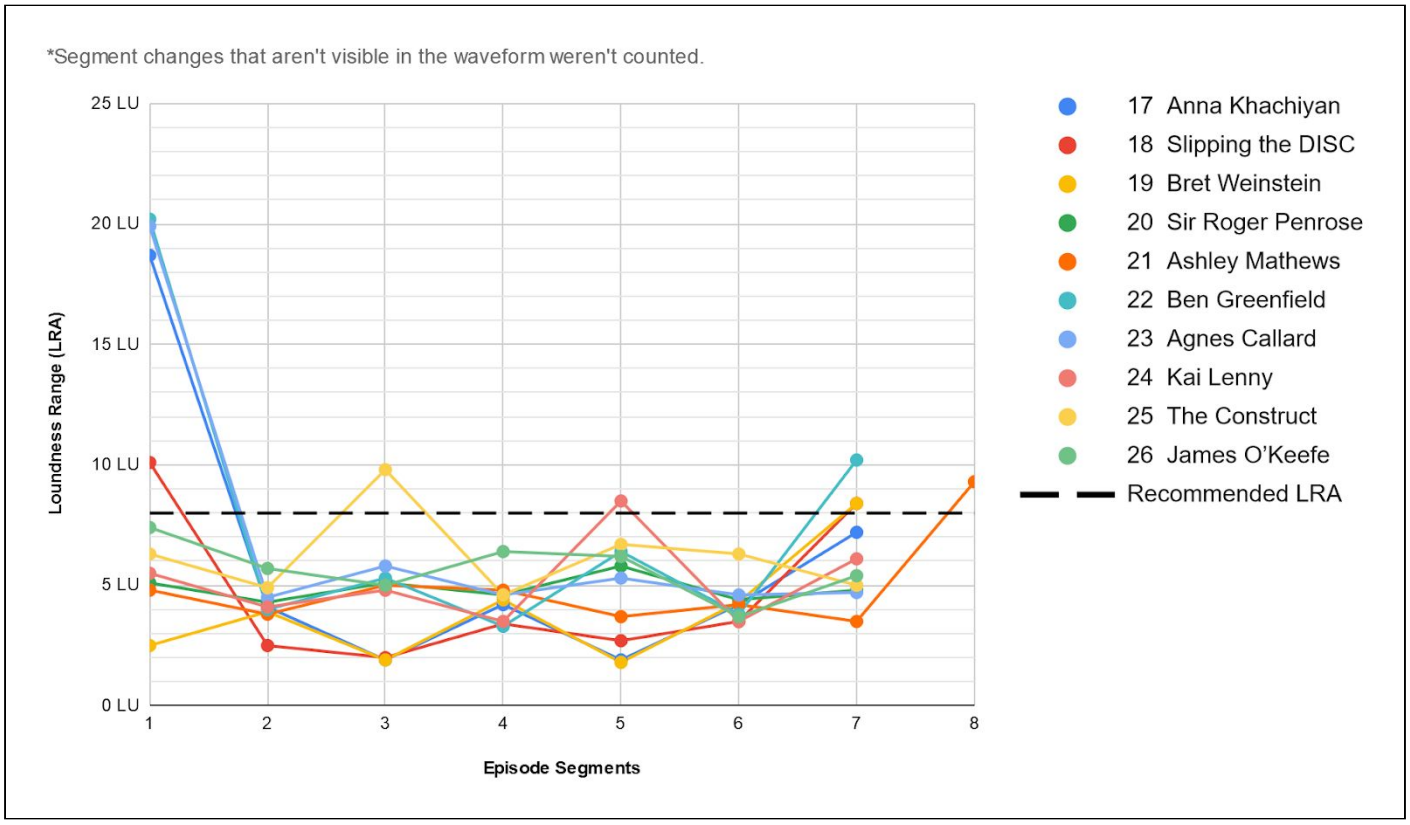

Seven of the last ten episodes break the recommended LRA at least once. Some offenders are more serious than others, but I'm not worried about them at the moment. We've got bigger fish to fry.

Consistency

This is where things go off the rails. Segments within an episode and between episodes are wildly inconsistent. As a result, episodes are inconsistent and therefore the podcast as a whole is inconsistent.

The graph below shows the perceived loudness (LUFS) of each segment per episode for the last ten episodes. Keep in mind that the goal is -16 LUFS to -20 LUFS. There's some wiggle room around those standards, but not enough to tolerate the variation below.

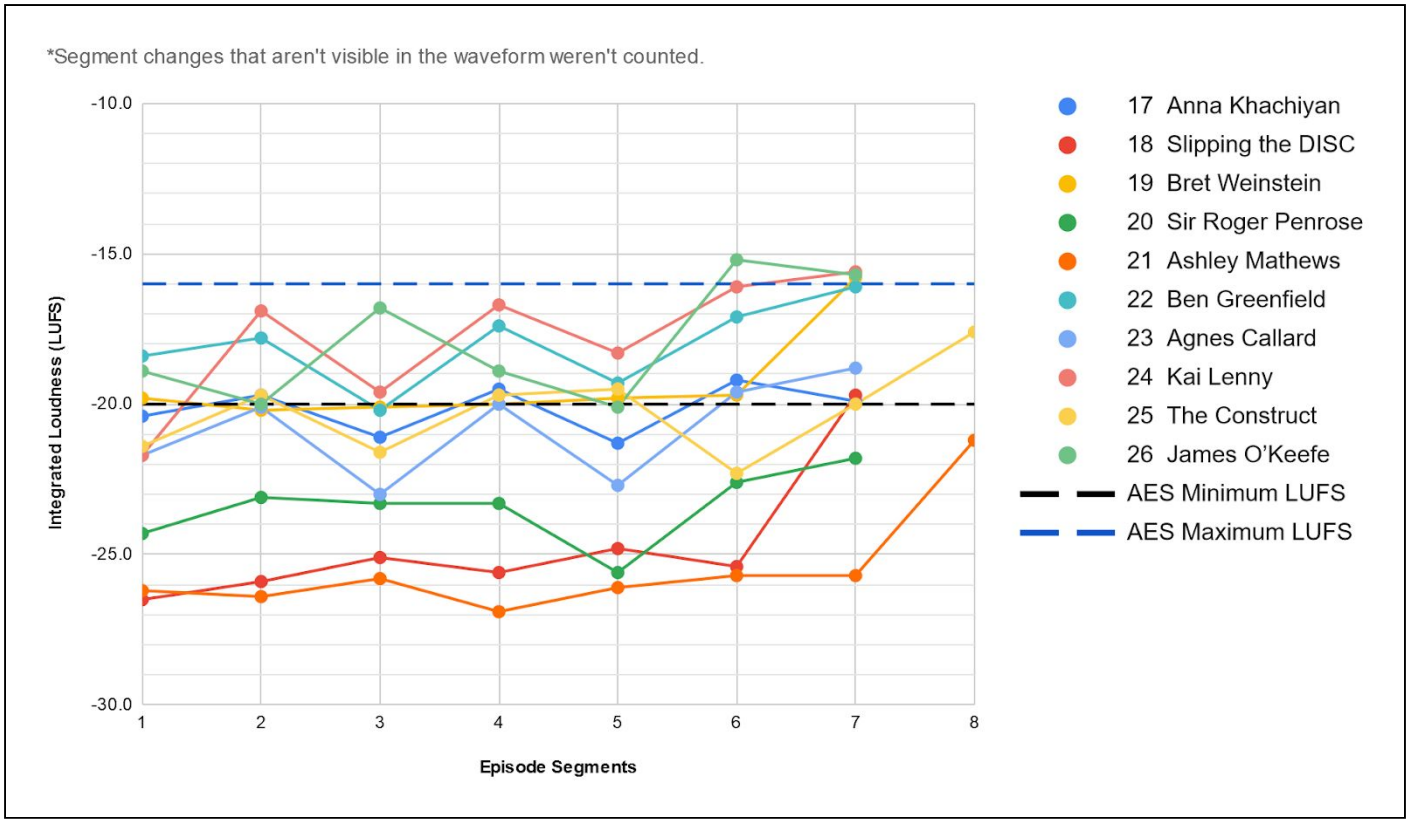

Now we have a problem. Some listeners will need to make repeated adjustments to their device. Others will need to switch to something more powerful, like their car stereo, home stereo, or computer monitors (depending on type). If the unwanted noises mentioned earlier haven't been removed, they'll have to contend with them too. Or they could forget about the podcast and spend their attention elsewhere.

Can this be fixed?

Yes, absolutely. But I see no signs that the above problems are being solved. It's clear that some of the audio is well processed. Unfortunately, that quality isn't consistent between segments of an episode or between episodes.

Without knowing more about Eric's production and post-production processes, I can't lay blame in any particular direction. The best I can come up with is highlighting some questions worth answering:

- Why is there such inconsistency within episodes and between episodes?
- Can the current post-production process be modified to improve audio quality and overall consistency?
- If it can't, can someone or a group take over the responsibility of audio post production?

I'm happy to help out with the latter.