

StoryCo as a World Model.

The show is already an intelligence. A proposal to point our world-model engine at StoryCo - and turn a run of long-form conversations about the Story Economy into a queryable model that the show, and its guests, own.

Reads alongside *One Engine, Many Networks* (decent-rw.zocomputer.io/articles/one-engine-many-networks), the outward-facing explainer of the engine described here. One document is what the engine *does*; this one is what it would do for you.

01 The idea

Every week StoryCo puts a thoughtful person in a room and has a long, real conversation with them about how story creates value. That is not just content. It is exactly the raw material the engine is built to read: attributed talk, over time, from people who think about this for a living.

Point the engine at the show and you get a *world model of the Story Economy* - a continuously refreshed, queryable picture of what your guests collectively know and keep arguing about. Not a transcript archive you scroll, but an intelligence you can ask. The show stops being only a feed of episodes and becomes a body of knowledge that compounds with every guest.

02 Why StoryCo, specifically

The premise of the show is the premise of the model. You are already gathering, week on week, the most articulate people in and around the business of story. The thing that usually gets thrown away - the reasoning, the half-formed frames, the rhymes between one guest and the next - is the most valuable part. It is the same insight the AI labs stumbled into: the exhaust turned out to be the engine.

So the question becomes: who owns the intelligence that accrues as the show grows? Right now it disperses into the feed. Modelled, it becomes an asset that belongs to StoryCo and the guests who built it.

03 Why it's an interesting experiment

For us this is a genuinely new shape to prove. Every instance we have run so far reads either a chat room or a weekly call - private rooms, text-first. A published podcast is a different kind of media to ingest:

- **Long-form and spoken** - argument, not chatter. Denser, more considered, harder.
- **Multi-source** - the audio/video, the transcript, and (optionally) any production notes are several views of one conversation. The engine has to compose them into one model.
- **Public-facing** - a network you publish *to the world*, not a private room you control.

In the protocol we describe four shapes we want one engine to survive. A multi-source creator network - public video plus podcasts - is the one we most want to test, and StoryCo is the cleanest version of it we have come across. If the engine reads StoryCo well, it reads almost anything.

04 How the process would work

The same five-stage engine, with one new adapter on top:

1. **Source** - point an adapter at the episode feed and transcripts (and any production chat you choose to include). It writes each episode into a uniform stream of who-said-what.
2. **Normalise** - each guest becomes one durable record; each episode a session.
3. **Extract** - per episode, the model pulls the frame, the key quotes, the open questions, the decisions, and the concepts - attributed and dated.
4. **Distribute** - concepts that recur across guests promote themselves; edges form between people and ideas; a cross-episode graph builds.
5. **Publish** - a queryable StoryCo world model, a page per guest, and a short digest after each episode: the diff of what the show learned this week.

We ingest naively - voice, timestamp, raw words and provenance only - and re-derive the meaning ourselves, rather than trusting anyone's tags or framing.

05 What you'd get

- A living map of the show's intelligence - ask it a question, get answers grounded in what guests actually said.
- A facet per guest - their spoken corpus, their themes, their connections to other guests. Something each of them would want to share.

- Recurring themes surfaced across the whole run - where, say, Will Storr and a brand strategist and an economist are unknowingly saying the same thing.
- A digest after each episode you can publish to your own list.
- And, in time, a layer the show owns - the direction of travel for anyone serious about not handing their value to the platforms.

06 A small first step

Low lift, no commitment. Take the episodes already in the can - including mine - run the engine over the transcripts, and produce a first cross-episode snapshot for us to look at together. A week or two of work, and we would be reacting to something real rather than to a slide.

07 Whose record it is

One principle underneath all of it: the person signs for their own record, not the host. What a guest makes public and what they hold back is set at the boundary, by them. The model reads the room only on terms the room agrees to. That sits in our consent spec, which travels with the engine.

The aim isn't to archive the show. It is to let StoryCo keep the intelligence it is already generating - and own it.