Let’s say you walk around a corner one day, and you see a guy kicking a dog. That guy is an asshole, right?
But what if you find out that the guy kicked the dog because it was attacking him before you came around that corner?
But what if then you found out the dog was attacking him because it went nuts – because it ate a bunch of crack that fell out of the guy’s pocket?
But what if we then find out he bought those drugs to give to his dying wife to help ease her pain because she doesn’t have medical insurance?
But what if we find out he’s gaslighting his wife, and she’s not really sick at all – he’s tricking her into thinking she’s dying because it makes him feel important.
We’ll come back to this later.

BEHIND THE SPHERICAL STAGE

TAKING VR STORYTELLING BEYOND GAMES AND MOVIES

FECHA / DATE: VRDC 2017

TÍTULO / TITLE: Behind the Spherical Stage

TEQUILA WORKS

C/ Serrano 19, 5º izquierda, 28001 / Madrid (Spain)

www.tequilaworks.com

Good morning everyone.

We are Rob Yescombe and Raúl Rubio.

We have worked together on RiME and a VR experience TBA soon.

There's a scene in the Scorsese movie THE DEPARTED where Jack Nicholson's character – he's a gangster – he shoots a woman in the back of the head. And when she flops over onto the ground he says to the guy next to him "Geez. She fell funny". Then the scene ends.
That line – that disconnected reaction to her death – it says a lot about his character. But what a lot of people don't know is that when Scorsese shot that scene, he kept the camera rolling - and Jack Nicholson kept improvising after that line. So at first, he said "Geez. She fell funny." But then he paused looking at the dead body and he says "Wait. I think I want to fuck her again." Then he paused, and he said "Ahhh, I'm just kidding". It was just a joke.
In the final edit, Scorsese could have ended the scene after any of those three different lines, and each possible version would've changed our perspective on who that character is – we weren't ever going to see him as a nice guy, but it means the difference between us seeing him as a guy with a very dark sense of humour versus a guy who literally has sex with dead bodies.

In a linear story like a movie or a book, we only ever see <u>one perspective</u> – the perspective of the author. We're only ever allowed one conclusion.
We will only ever see a character in the way the author intended. The reader is not a collaborator in the experience of the story.

One of the best things about games, is that you and I can have different experiences of the same moment.

These kinds of *contrasting experiences* are great at generating word-of-mouth for your project. People want to share and compare.

This kind of thing is easier to achieve in games where the focus is on ACTION; where there are lots of variables that can be generated by physics and mechanics.

But it’s harder to create this kind of *organic-feeling* contrasting experience in a game’s STORY.

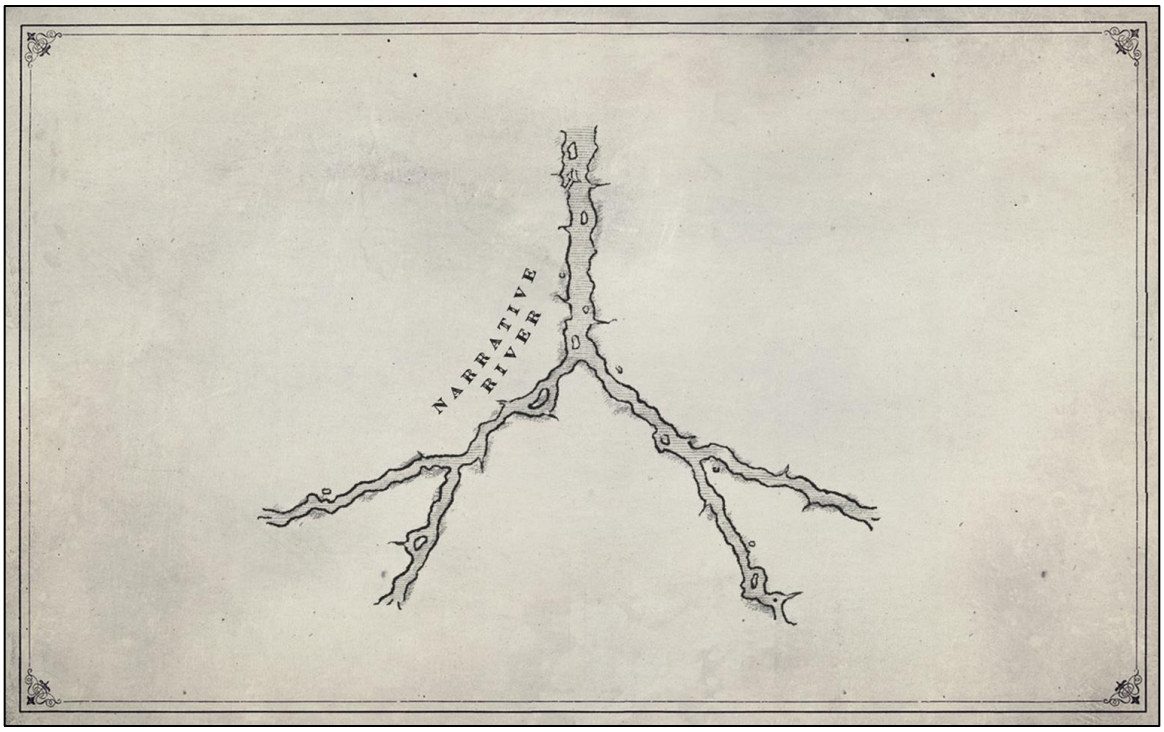

In the effort to create *contrasting experiences* in a game's story, the common solution is to use a branching narrative: choose the left door or the right door; be nice to a character or be nasty to them; save someone or let them die.

This works pretty well. It creates variables between players. And I like lots of games that use this technique - but it never feels like real life; it feels mechanical.

At first we thought it felt unnatural and mechanical because the moments of choice are so clearly defined: you always know it's time to make one of those choices, because two buttons start floating in the air – or something like that - and I have to press one to continue.

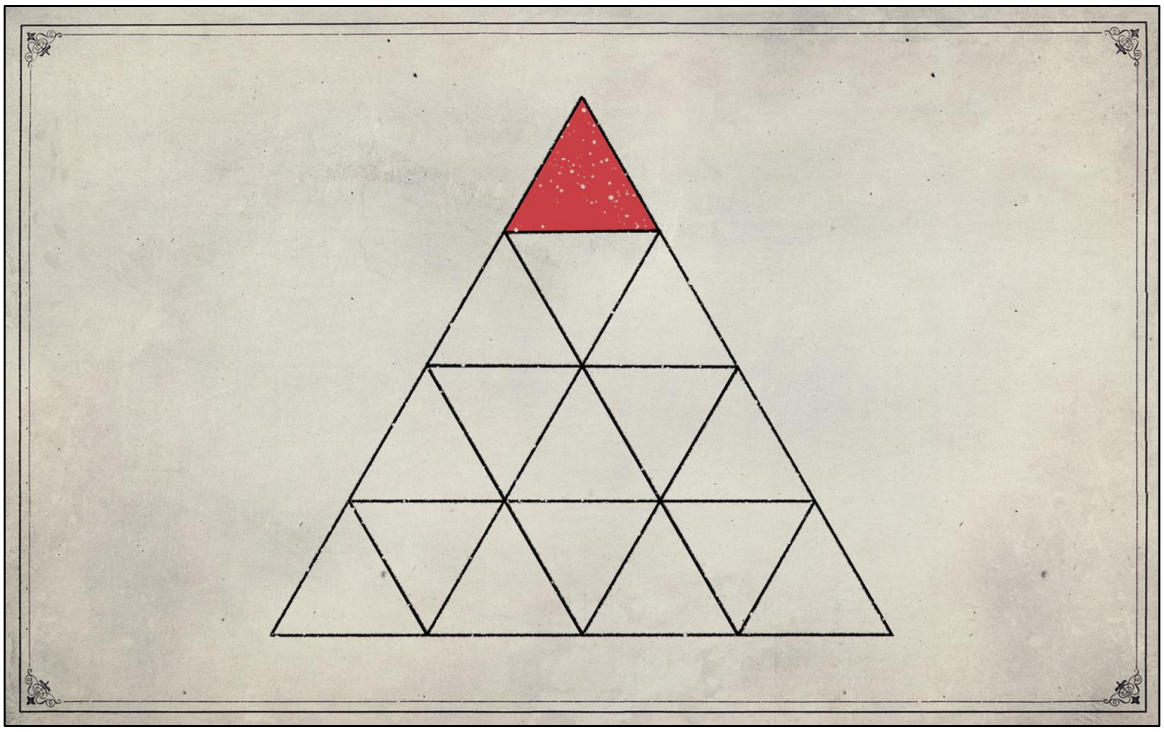

But on reflection, I realised that it’s actually because those games place me at the top of a narrative pyramid – where I’m the most important thing in the world. Nothing matters until I show up.

Sometimes, that feels great – sometimes I want to be the most important thing in the world. But in our pursuit of Naturalism, it doesn’t feel right – it doesn’t feel REAL. The pyramid may not break my immersion, but it does bend it.

Because what I’m looking for is *a world that seems like it would go on existing without me*.

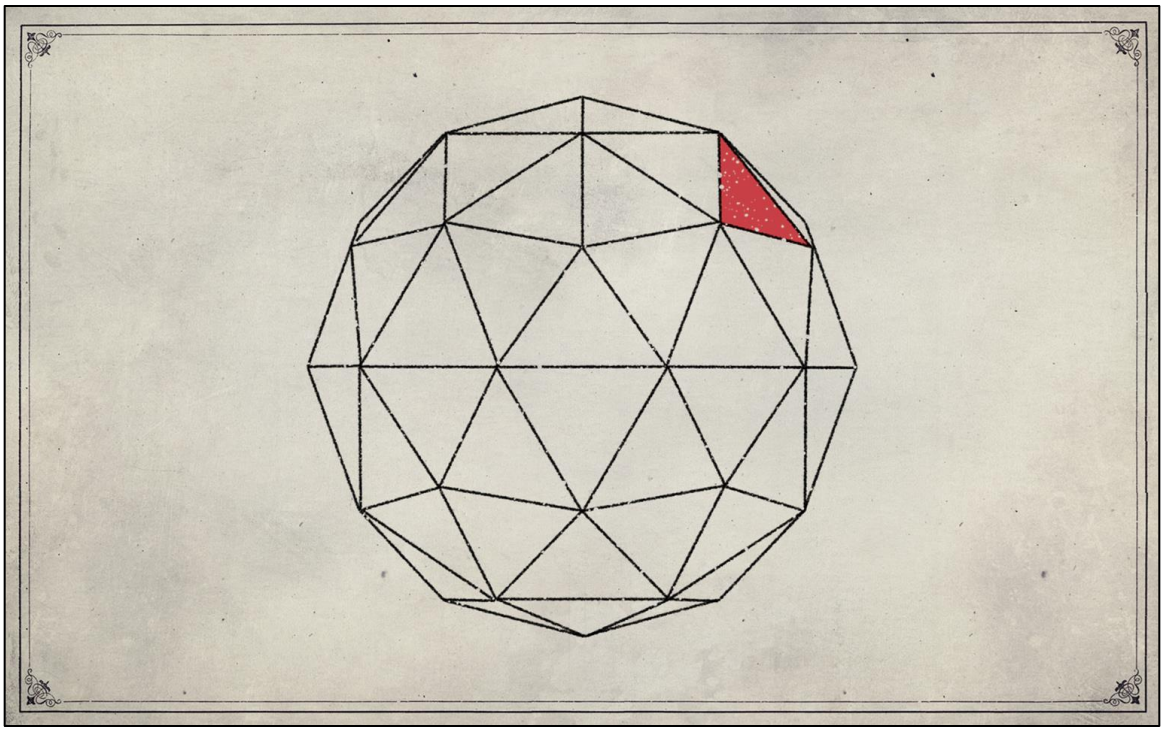

Here’s the weird thing about real life:
Right now, I am the hero of my story. And [YOU] are just a supporting character in my story. But at the exact same moment in time, [YOU] are the hero of your story – and I am just a supporting character in it. And all these people – they’re just background characters in both of our stories.
But at the exact same moment in time, each of them is the hero of their own story – and we are both just background characters in them.
That’s the narrative structure of real life.
Real life isn’t the pyramid.
It’s a sphere. Where we all share equal importance at the exact same time.
But here’s where the experience of real life really differs from entertainment media:
Our opinion of people and events changes depending on when we started looking. Two people can have completely different perspectives on the same person.

Think back to that guy kicked the dog. I told you all his secrets – the good and the bad. But in life, in the sphere, we don't always see everything in each other's lives.
Let's say you only saw the part of his story where he was taking care of his dying wife. Well then you'll believe he's a good man.
But if you only see the part where he kicked the dog, well then you'll believe he's bad.
But what if you only saw him gaslighting his wife – well then he's the God damn worst.
And the degree to which you think he's a good or bad guy has tons of possible variables based on the combination of good and bad behaviour that you may or may not observe.
And that got me wondering: what if we structured a story in VR in the same way real life is structured?

Imagine yourself in a dark theatre.
Imagine you're there to watch a murder-mystery play. And on the stage is a huge Agatha Christie style mansion – split right down the middle and opened up, so you can see all the floors and all the rooms.
Now imagine two characters in one of those rooms, discussing who they think the killer might be.
At the end of that scene, one of those characters will go down to the cellar. The other will sneak up to the attic.
But now imagine getting up out of your seat and walking onto the stage; into the play - into that dusty old mansion and following one of those characters to see what they do next.
But here's the thing: whoever you follow, the other character's story will still be happening at *the exact same time* in the other location. You could even run back and forth between them if you want.
Simultaneous stories.
But now imagine this play has eight characters.
Eight unique stories. All interconnected. All interdependent. All simultaneous. *Just like real life.*
You can move freely between all of them, any time you choose.

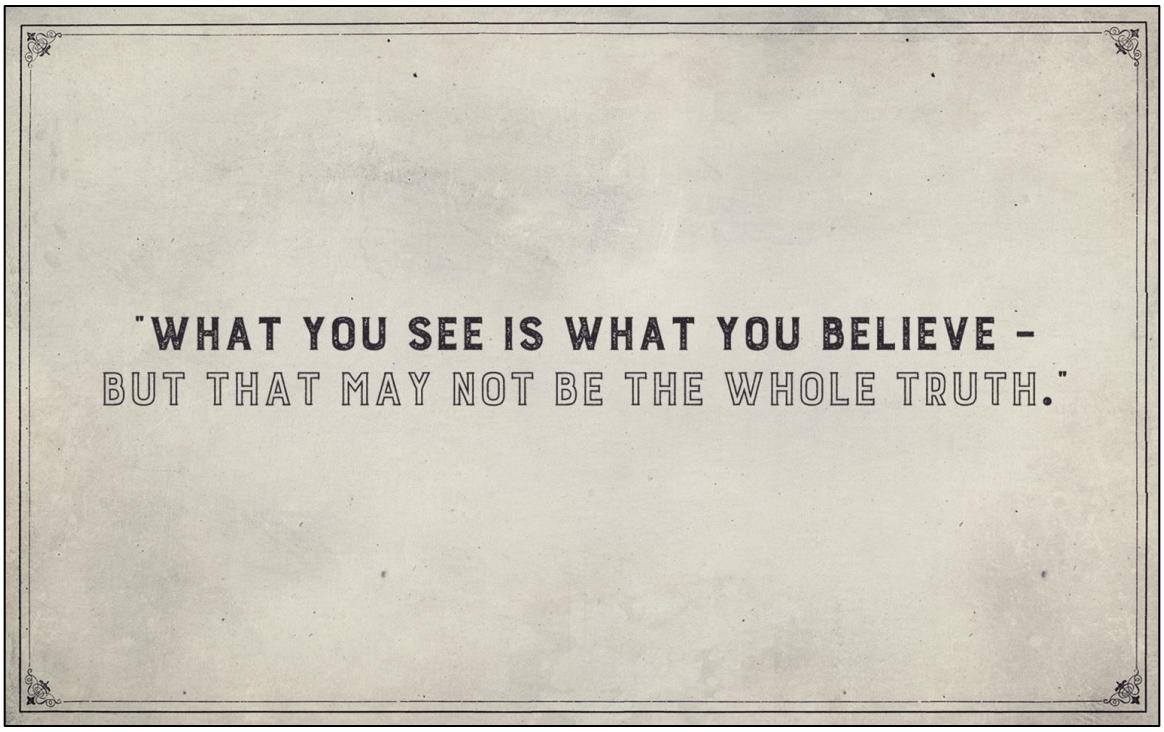

Now, although we're not allowed to say by name yet, *that* is the nature of the project I'm making with Tequila Works in Madrid: an eight-character murder mystery that is one hour long – but within that single hour, eight hours of story is taking place *at the same time*. The challenge to you is to decide how to follow it all.

And at the core of that experience, is this statement:

What you see is what you believe – but that may not be the whole truth.

This is our first attempt at this kind of narrative experiment, so right now you're not a character in this world. You don't change the story - *but* your story experience changes *completely* depending on who you follow and when. That means that you and I could reach different conclusions about who the killer is – just by virtue of how we observed the exact same content in different ways.

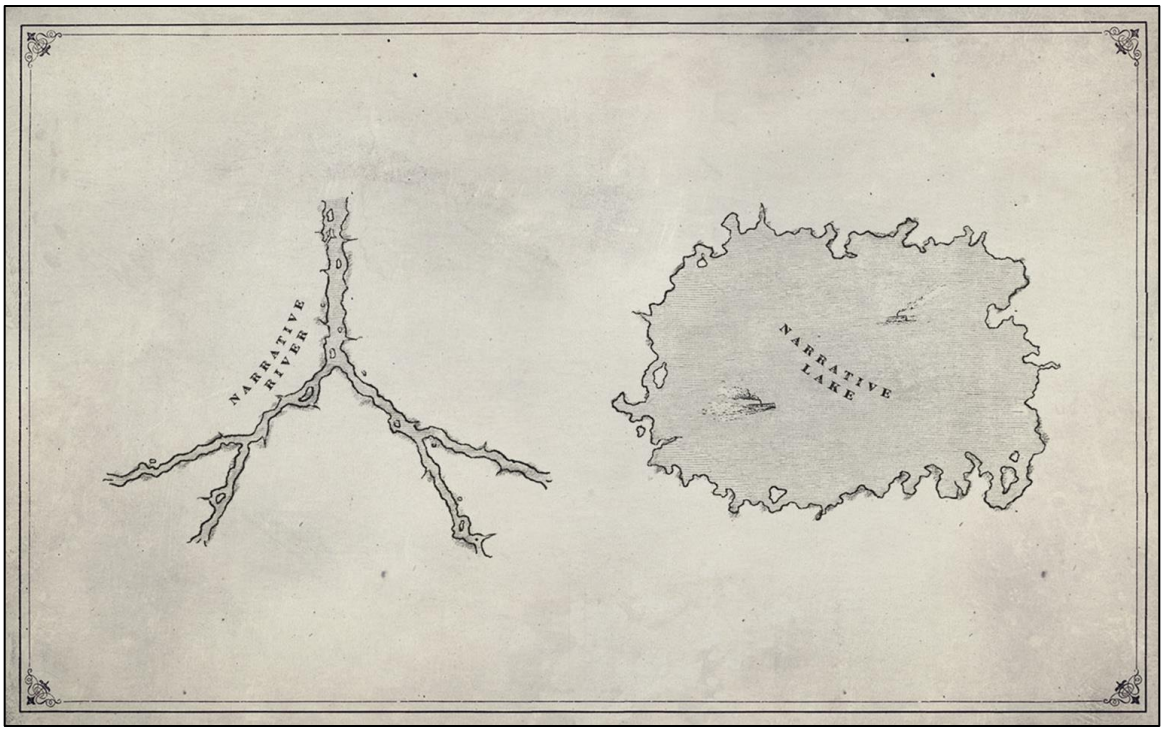

So, rather than a series of binary choices, instead you're presented with a single unit of time that you can explore any way you want. Freeform and organic. It is the feeling of a narrative sandbox.
But here's where it gets really interesting: let's say you follow the Detective from the very beginning – well then the story will look like a detective story. But if you suddenly decide to look the other way, and follow someone else for a while, you might start to think you're in a love story not a detective story. But then if you follow someone else, you might start to think you're in a science fiction story, or maybe a horror story, or even a comedy.
So not only will your understanding of the core mystery change depending on how you explore; your perception of the actual genre of the story itself will change completely organically – based purely on who and what you're looking at, at any given time.
And, despite how complex this all sounds on paper, this is the structure of real life. So, whilst it's challenging for us to write it and make all those pieces fit together, it feels incredibly natural when you're inside it – even as an observer.
It makes a story in Virtual Reality feel like it's virtually reality.
Now even though you could technically make this as a regular game, we know that the simple act of *looking at things* is one of the most compelling things about Virtual Reality – so at its heart, this is a kind of game of observation. Which means, you can take things at face value, or you can explore and dig deeper into each character's story. There's an opportunity for challenge – but only if you're in the mood to challenge yourself in any given moment.
It's also important to mention the distinction between this experience and a regular Walking Simulator – in most cases, those kinds of games put you in a very *lonely* place. You're usually exploring somewhere after all the action has happened, and you're piecing it all together with letters and audio logs, or watching ghosts play flashes of scenes; that kind of thing. But we wanted to create something that really felt ALIVE *while you were there* – and this structure unlocked that for us.

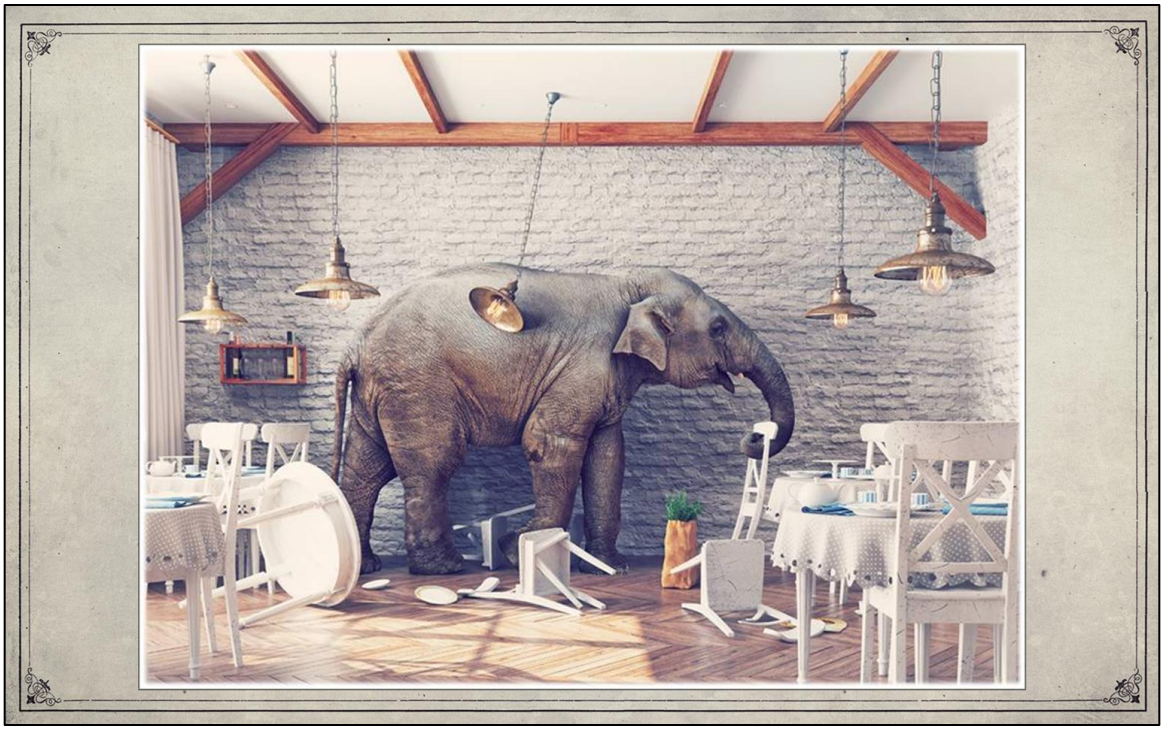

Now, I imagine some of you are thinking – if you can be anywhere in that house at any time, what happens if you immediately find the killer?

Well, first of all, whilst we market this product as a murder-mystery, you actually have EIGHT equally important stories happening simultaneously within that single hour. So if you 'solve' the murder very fast, and figure out who the killer is, there are still seven other mysteries happening in that house.

But also, murder-mysteries themselves actually come in TWO forms:

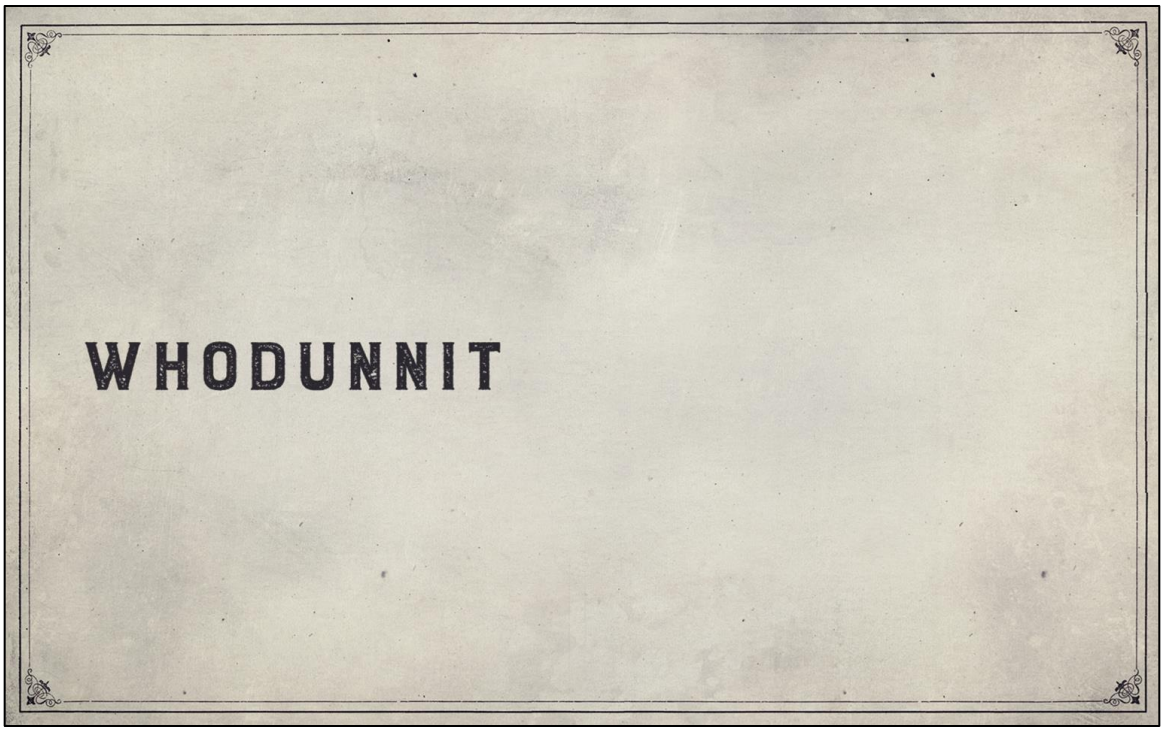

A Whodunnit covers things like Agatha Christie; the murderer is unknown, and that’s the puzzle.

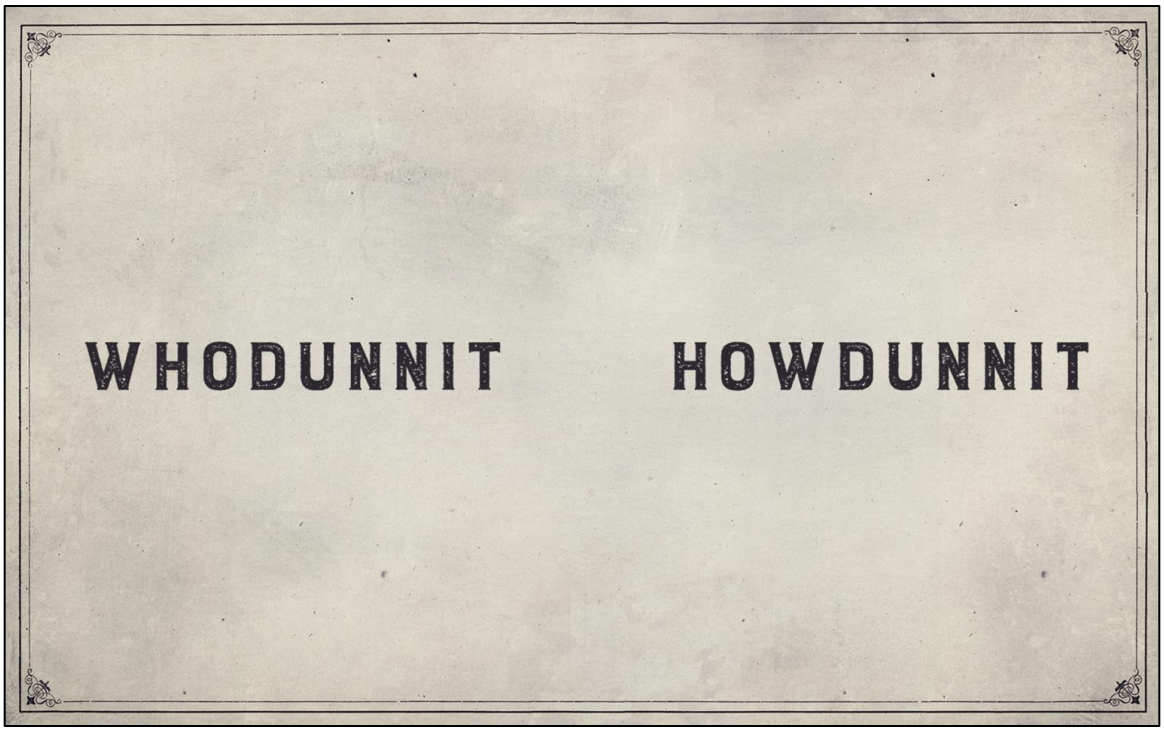

A Howdunnit covers things like Columbo or Dexter; the killer isn't an unknown - in fact they're often the star; we just want to know HOW they did it or how they’re going to GET AWAY with it.

Both of those formats can be equally compelling. So in the case of our project it can be a both a Whodunnit OR a Howdunnit, depending entirely on how you explore the hour.

Rob explained the theory. Now it’s time to put Spherical Narrative in practice.

We want Spherical Narrative to cover both **formats**: that meant creating an immersive real-time open narrative experience composed of multiple simultaneous storylines for each character’s point of view.

This spherical narrative format **presented** us with **unique** challenges.
To figure out we had to c**ombine techniques** not only from **games**, but also from **movies** and **theater**.

Today we’ll cover the 3 biggest ones.

1st challenge.

The most important thing to know about Spherical Narrative is that, technically, it’s impossible.
The whole concept is trapped by a paradox.

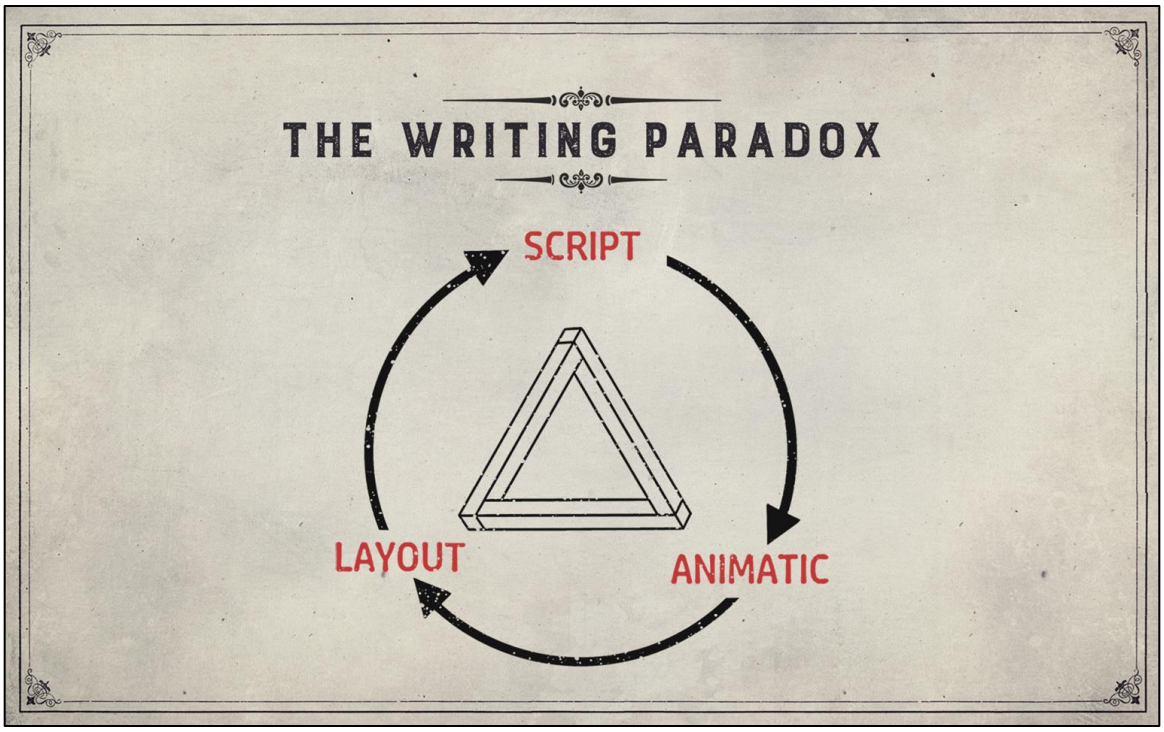

The script tells us how long every scene is –to the second-. for all the simultaneous stories to fit together
The layout tell us the space it all fits in.
The animatic tells us how these pieces fit together.
But the problem is, each one needs the others to exist firts.

<u>How do you know exactly how long all those scenes are until you write them?</u>
When you have several **interdependent** stories running at the same time, you can't just **sit down** and start typing.
You need a way to keep track of all the moving parts of the story.
An ANIMATIC is the **sensible** way to do that.
But you can't make that Animatic without knowing how long the scenes are.
But even if you could, you still can't make an ANIMATIC without a LAYOUT.
But in Spherical Narrative, you can't make a LAYOUT without the script – because when you're writing the script, you don't just have to factor in exactly how long each individual scene will be, you *also* need to factor in the distance between scene locations and the time it will take for a character to walk between them.

In games we usually rely on gameplay mechanics to define the scope and the experience.

The story defines time and space: it tells us how long will be the experience and how big will be the world.

In Spherical Narrative Story became the measure for other features and modules.

Suddenly the story was the most important thing in the project.

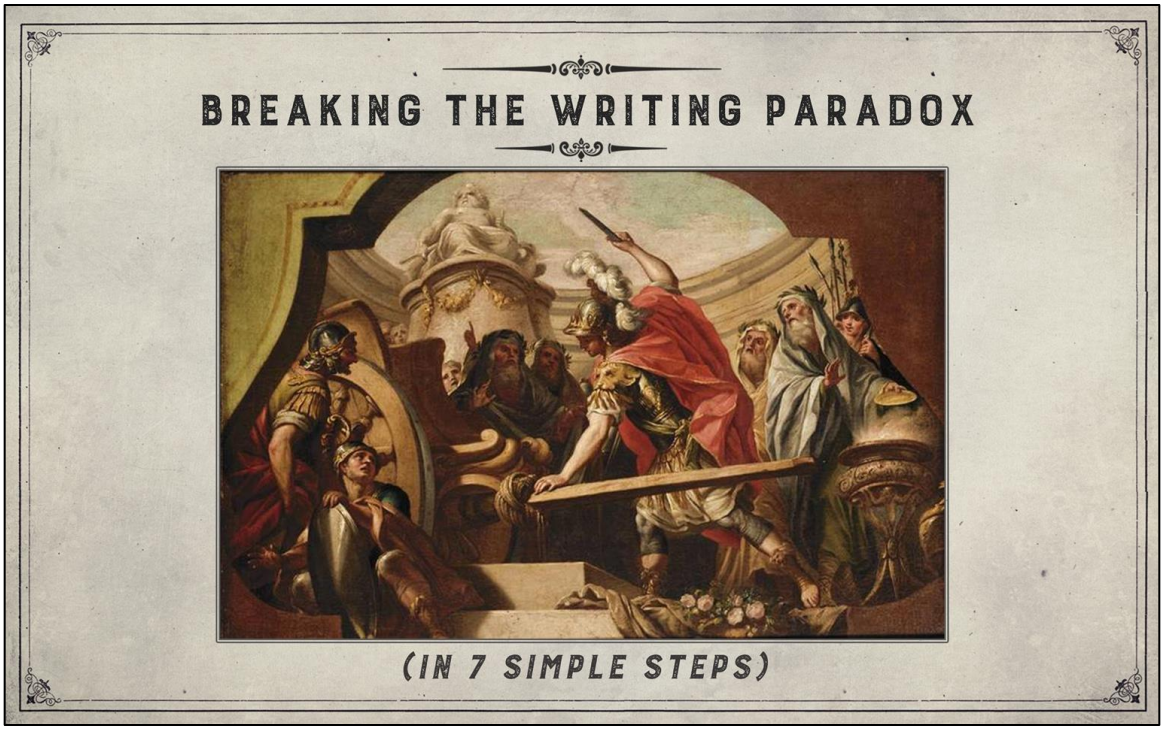

We had to take this leap of faith to break the Narrative Paradox.

With Story as our Sword we are gonna break it in 7 steps.

What's the experience?

Our focus in this experiment is immersion and curiosity. We needed to know how the story was going to awake the player's curiosity.

We want you feel like James Stewart in *the Rear Window*; you are a silent witness that can see everything.

There is no director, no editor of that information.

No predefined perspective or dramatic way to fake the truth. That's the player's role.

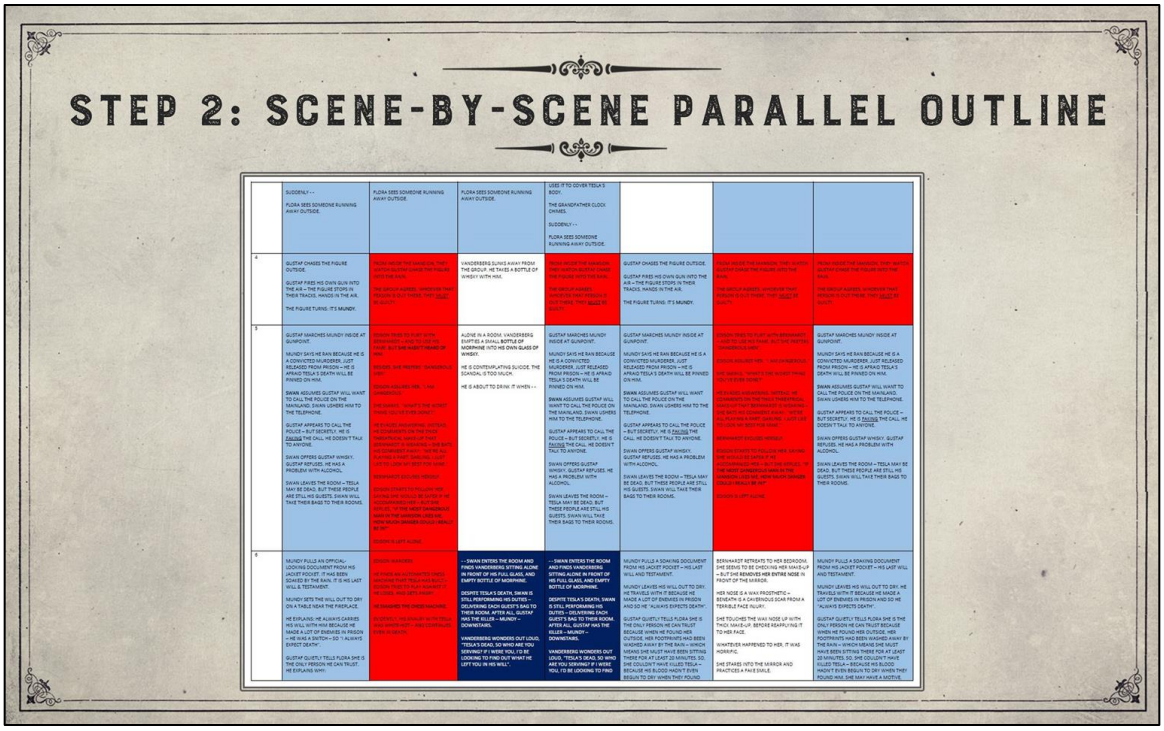

We created a scene-by-scene outline of each character’s story, where we could see them all stacked together in very approximate synchronicity.

A story with multiple points of view is hard to track.

The world is not there for you, so depending on where and when you are, what you are looking at, listening to or interacting with your perspective will be different.

And truth is a matter of perspective.

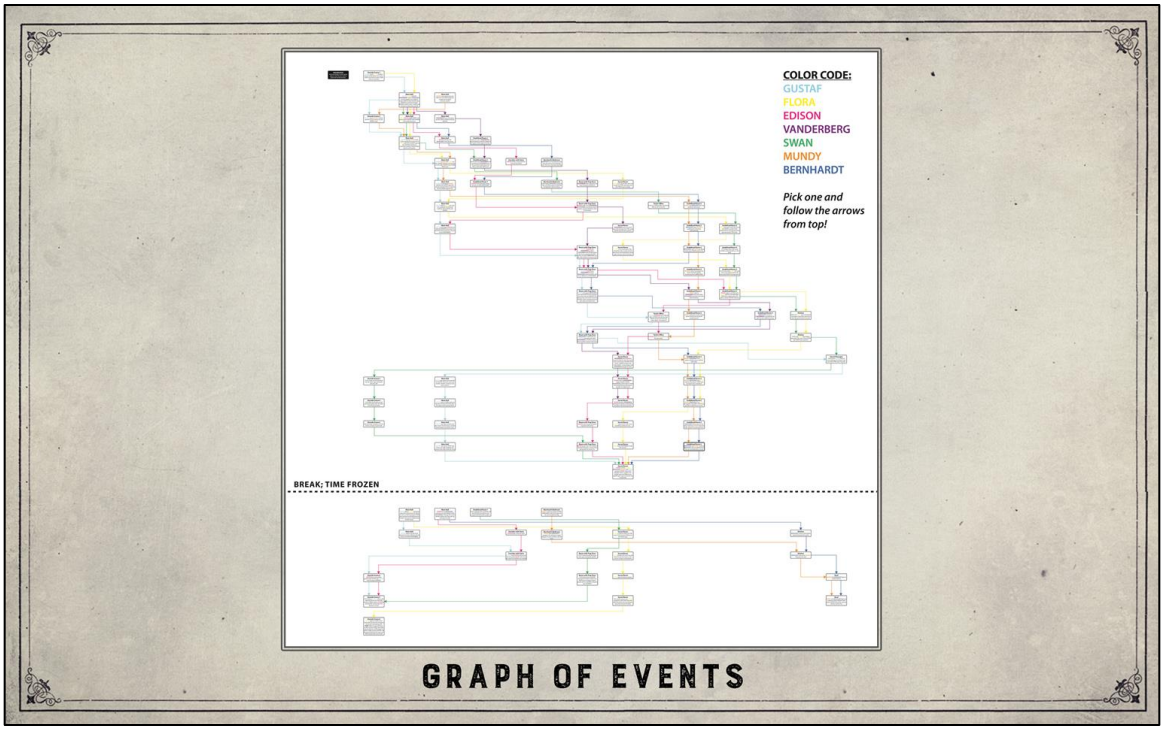

Then we created an easier to track graph that chronologically described the course of events.
Notice that the story has branches from the characters' points of view. BUT the **Player does NOT have branches**. Hence the narrative sandbox analogy Rob mentioned earlier.

This helped us to balance the timing of drama as **key events must happen at the same time in different locations**, creating exclusive information for the player, designed to change the perception and meaning of it for the player and complement (or contradict) other events.

Unlike real life, our actors could **not improvise new actions**.
But by splitting the world in **scene sets** and creating connections between them we had **safety buffers** where characters could idle, slow down or speed up to match the place and time they SHOULD be.

The graph allowed us to **adjust logic and avoid dead times where nothing really happened**, and allowed us to define those time buffers to readjust some actions i.e. not having 2 events next to each other.

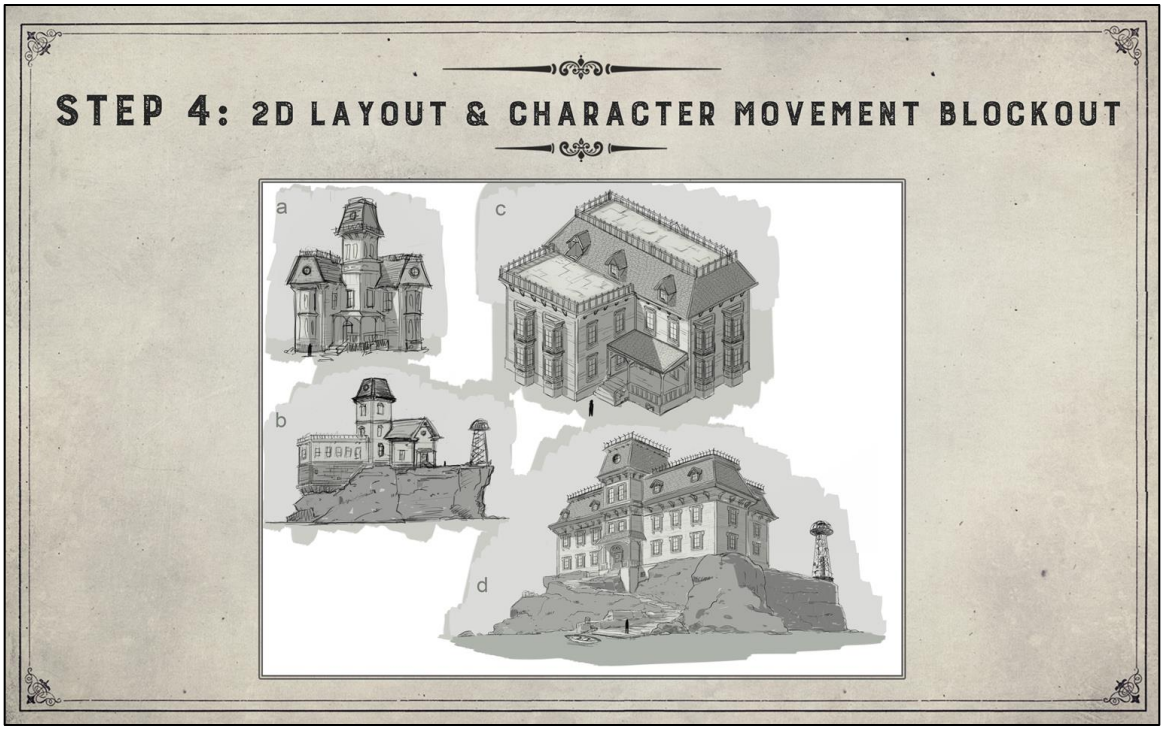

With a rough idea of the spatial requirements of the narrative –we could start planning an environment to fulfil those needs.

Because the story demands lots of characters sneaking around and deceiving each other, we knew that we were going to need a very flexible environment layout
– with multiple routes wherever possible –
but stylistically and architecturally **accurate**.

We had to trust that Rob could predict his own writing accurately enough to build the animatic around it.

<u>Being game devs we obviously prototyped our way to the answer</u>.

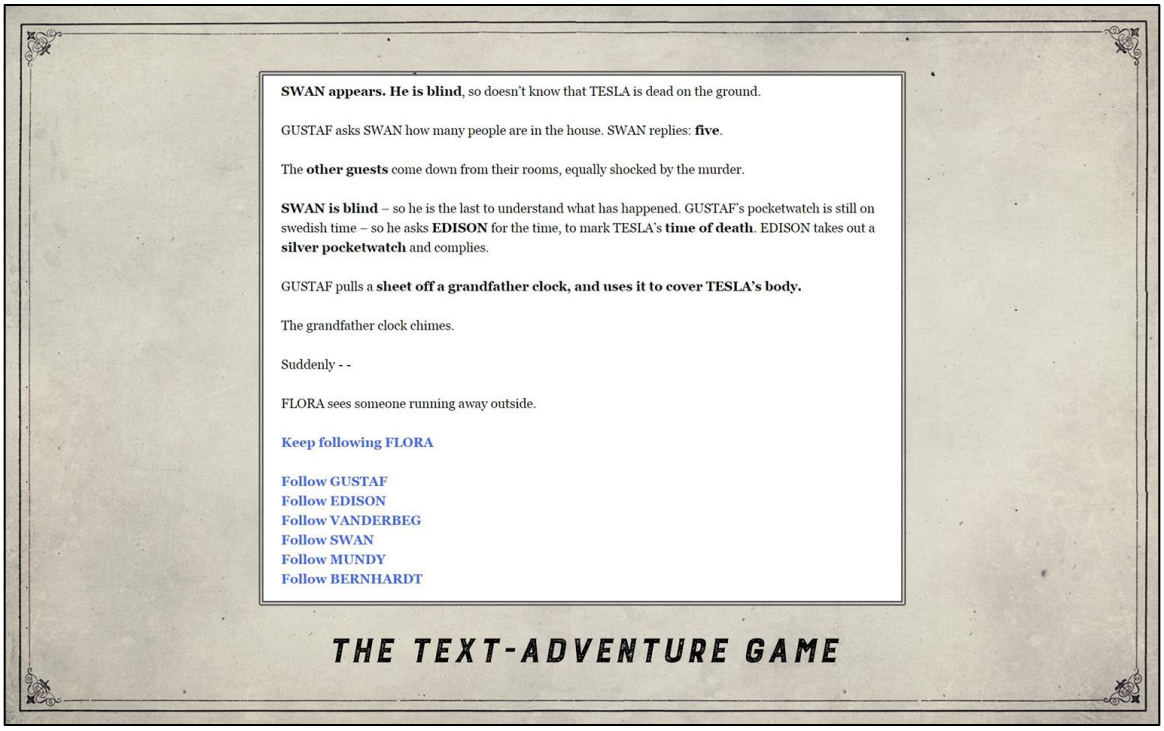

We prototyped a basic *choose your own adventure* text adventure and we play tested it.

Several more iterations later we had a closed enough narrative to draw a floor map.

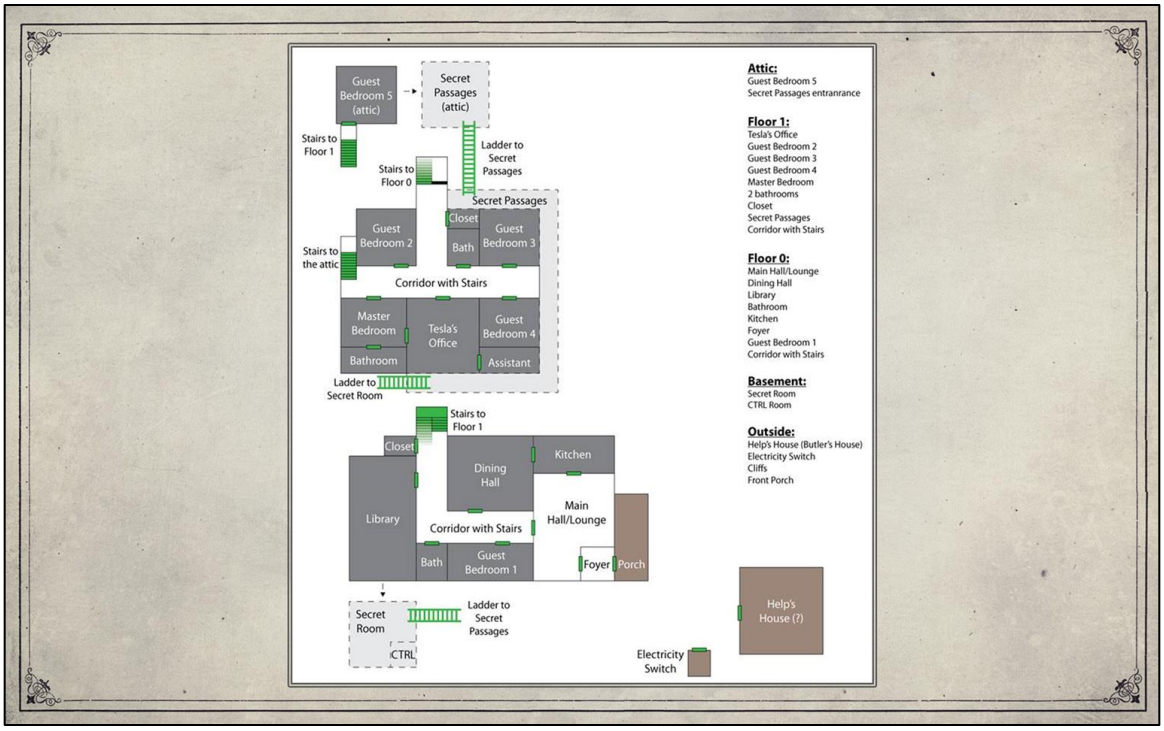

Using level design conventions, we made the 2D layout that incorporated all the story's key locations
–and also included additional rooms and paths where more flexible scenes could take place.-

Testing it against the outline, we chose where each scene would take place, each character's path and speed between those locations
–within a margin of around 10-15 seconds.

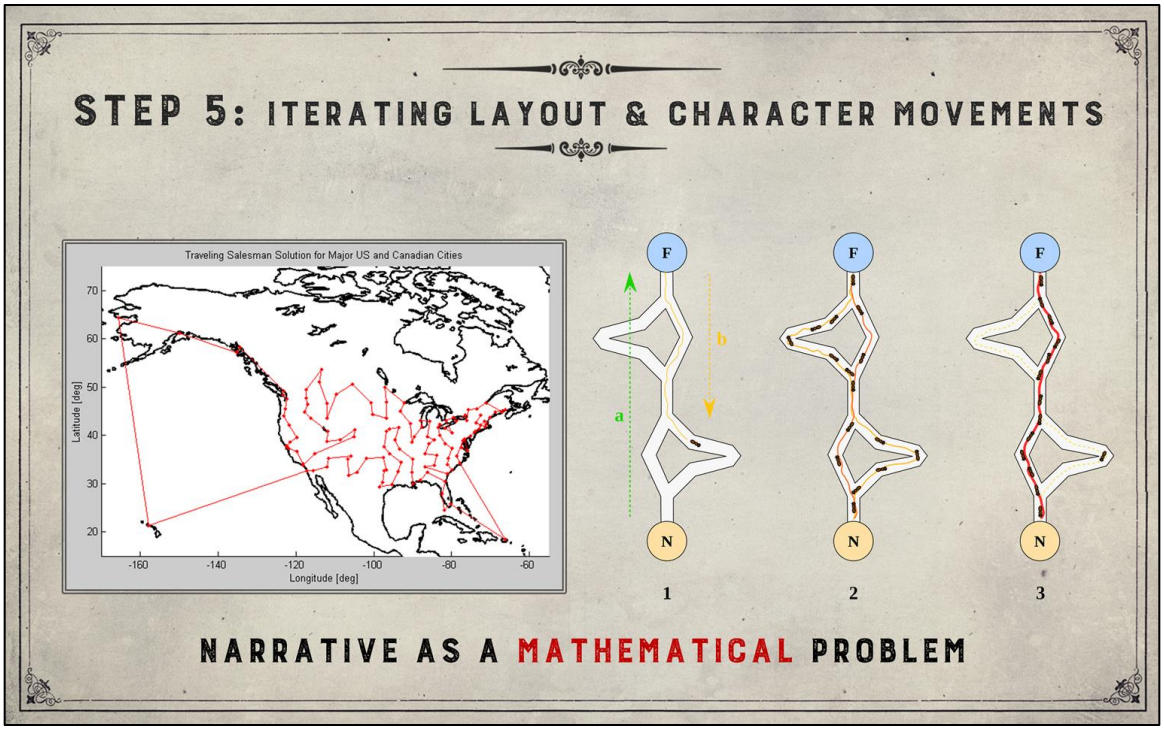

We then contrasted Narrative against Design, Art and Code, to create a layout that:

- reduced unnecessary movement
- without damaging the story and/or performance
- while being believable and interesting to explore

Optimizing the layout based on the story outline was not different to solving the Travelling salesman problem or the ant colony optimization.

Basically we had to optimize the route for a troupe of actors, between their respective sets to fully accomplish their respective scripts within very restrictive timeframes.

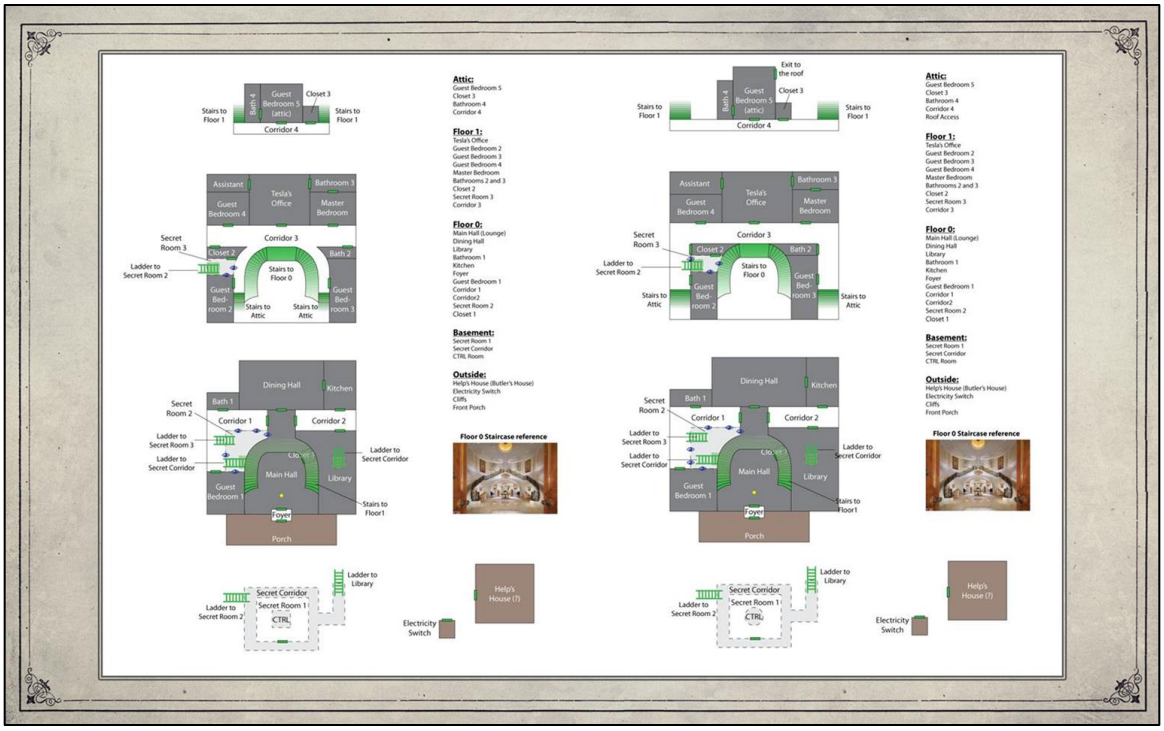

Actors either shared scenes together or monologues that excluded any other actors from that set at that time.

Not only that. To prevent any user discomfort or experience interruption, we avoided architectural details like spiral stairs, ladders, u-turns and elevators.

To travel between floors of the building, good old stairs became the solution.

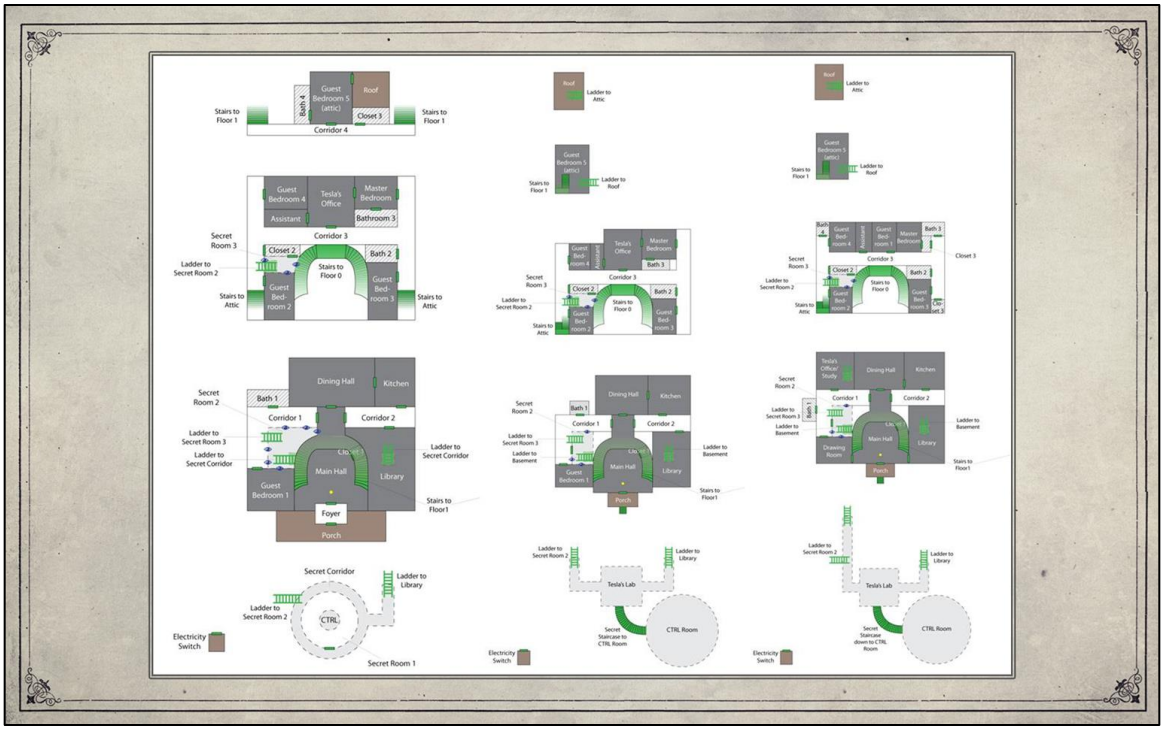

We went through around 30 iterations of this process until we finalized a 2D layout that worked for everyone.

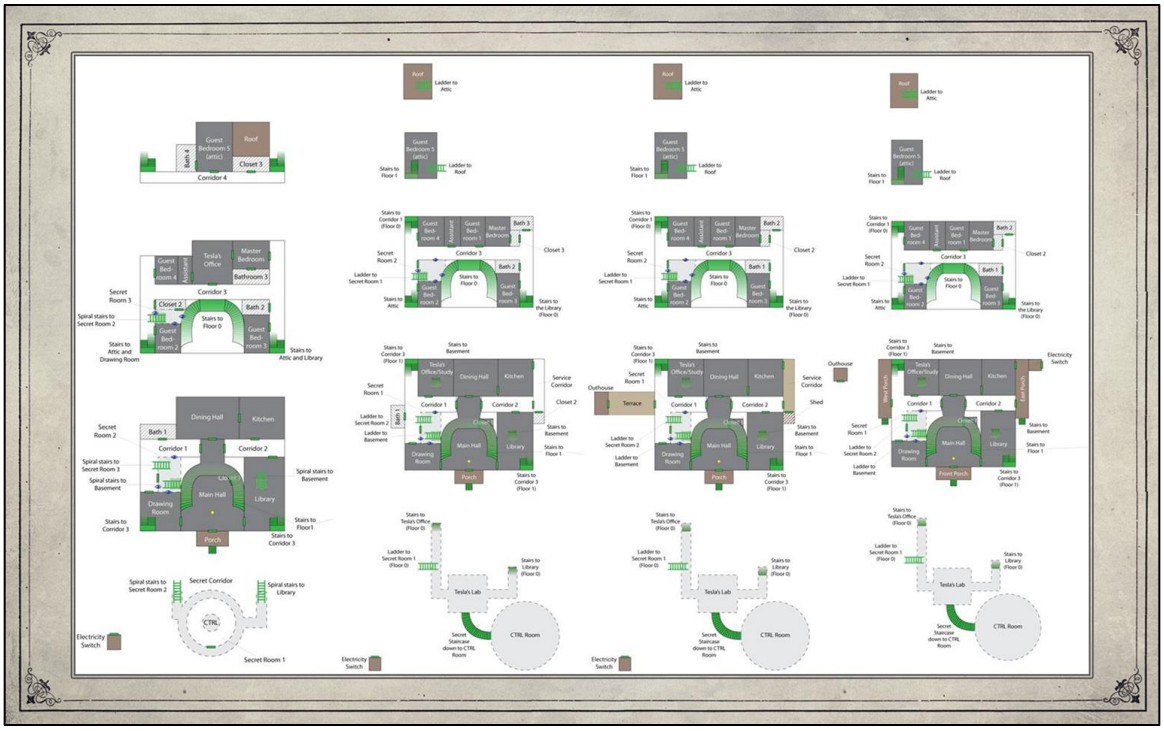

This is completely invisible to the player but the feeling was very different after applying smart level design decisions to the narrative.

It was a calculated risk – and perhaps if this was not an independent production, we might not have been allowed to take that risk.

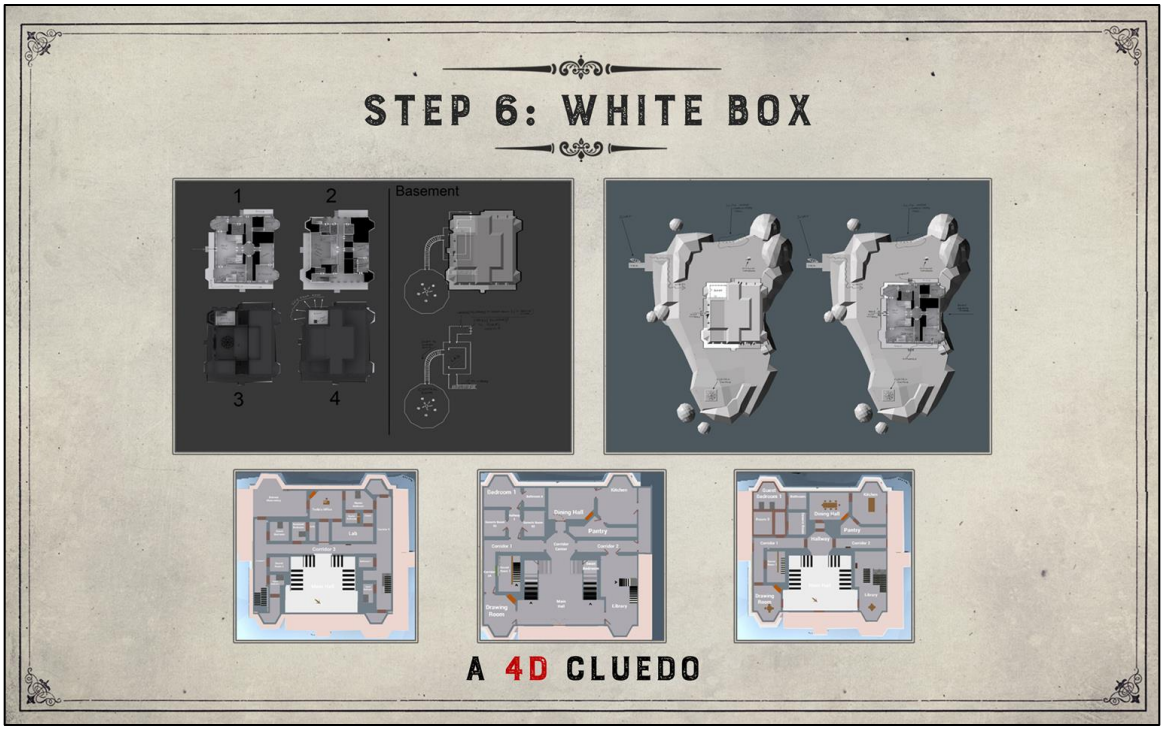

After all those iterations and revisions on paper we built a whitebox where we could test the character movements. Now we had some kind of 4D-Cluedo (Clue in the United States!).

Even in this primitive state we could see the potential problems it brought to Gameplay and level design.

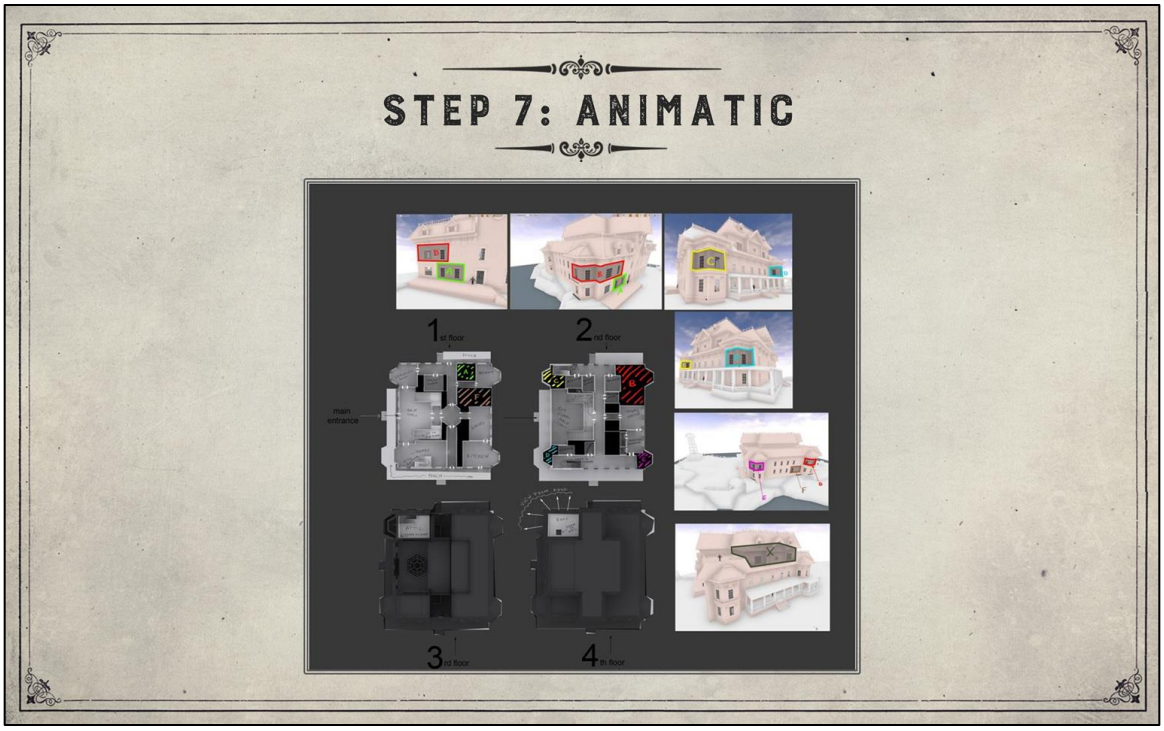

Next, we used the white box to plan out the exact movements of every character for every second of every story - roughly 8 hours of narrative – *all without a script*.

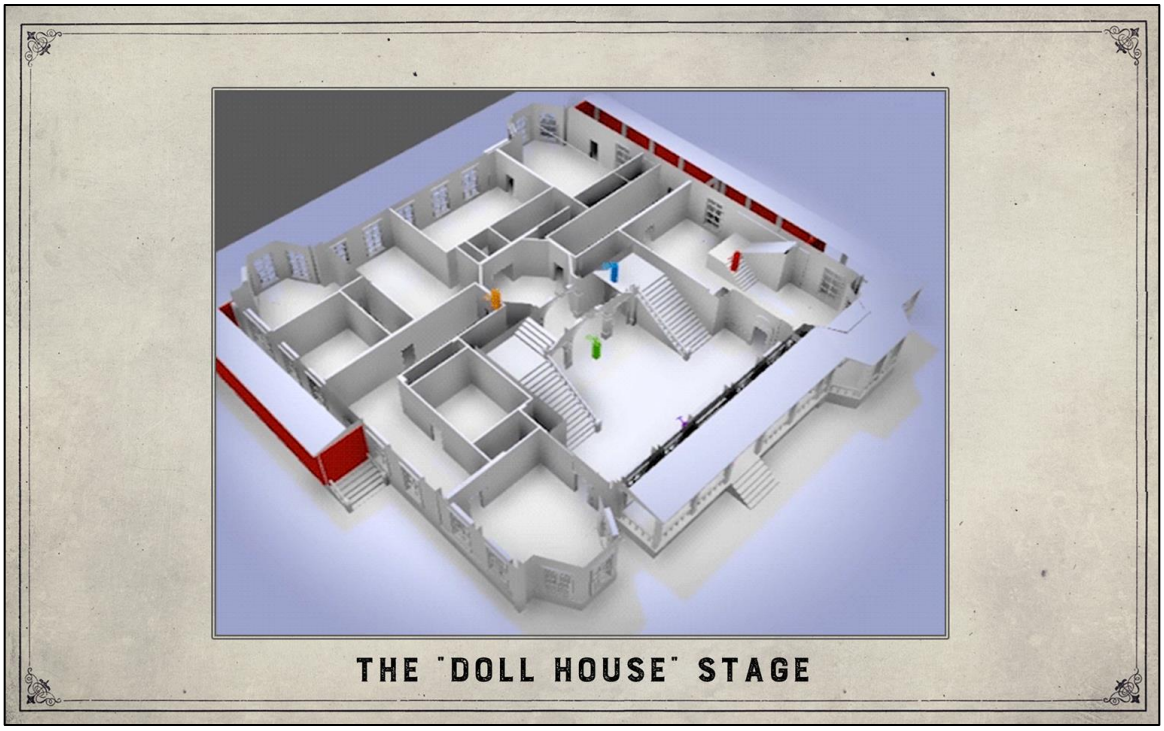

We iterated on adjusting the geometry and the character actions until we were sure the characters were not going through walls, corners or moved too fast.

That’s easier said than done but after a few attempts we had a proper animatic.

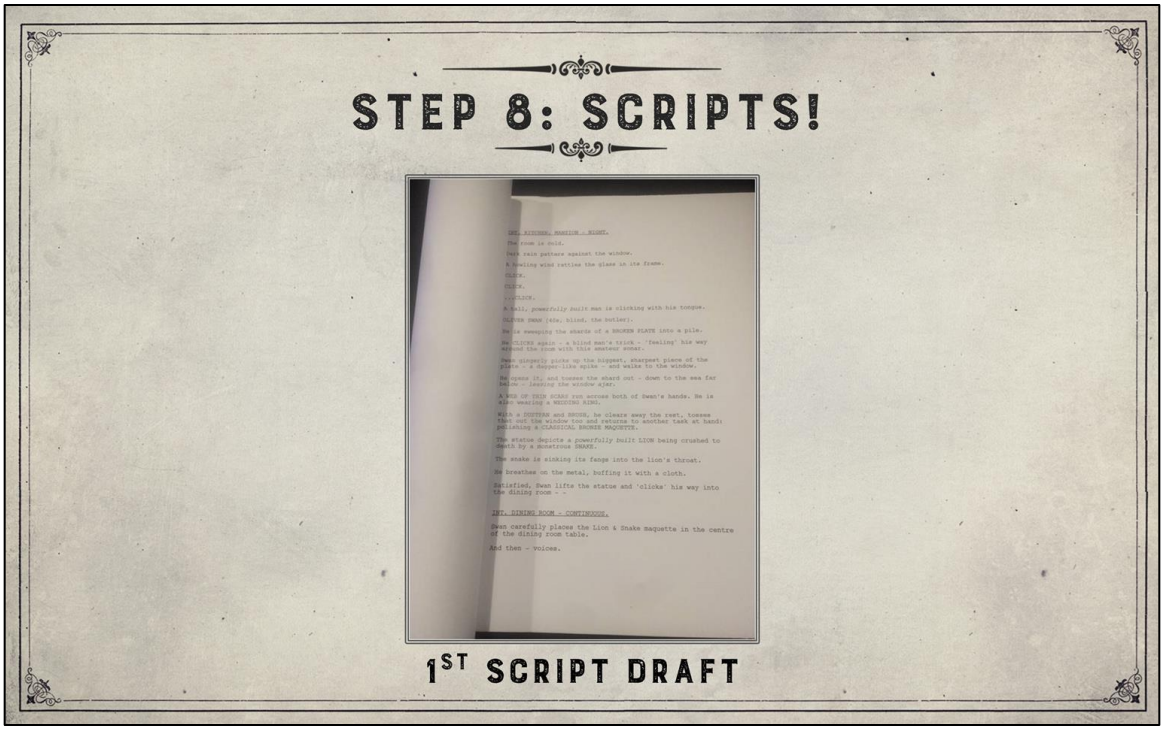

With the Animatic complete, Rob could finally move on the writing the full scripts – using the Animatic to make sure that his scenes matched each predicted scene length perfectly.

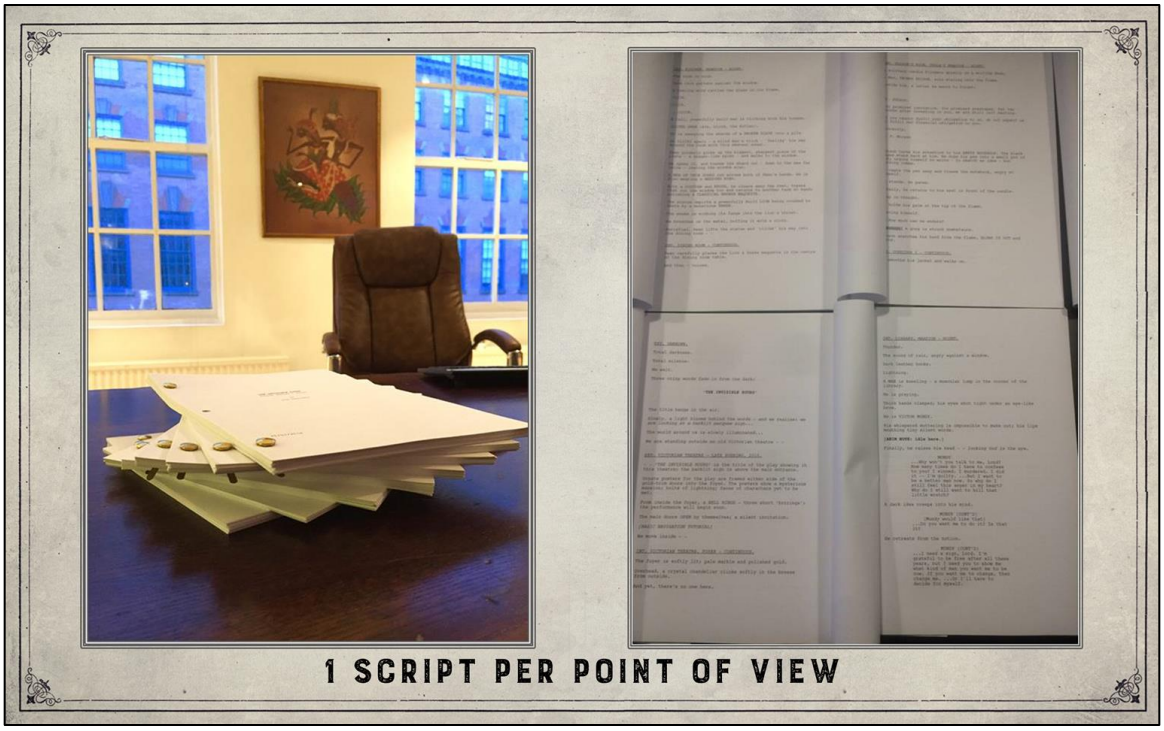

Finally Rob did customized scripts per character focused on their specific point of view.

Making narrative -an **unbreakable story-** the core helped us to define and prioritize technical questions.

A huge amount of that process was mathematical - but the truth is, there was simply no way to break this paradox without taking a leap of faith.

2nd challenge.

Immersion is the glue of VR.

Art can be astonishing and story can be awesome but are nothing without solid foundations.

Immersion is needed so nothing key is missing.

Immersion must be invisible to the player.

That means every scene depends on every other scene being the correct length of time.
If even one scene is incorrect by a single second, the whole system breaks.

If this was a piece of immersive theatre, and one actor was running late to a scene by a few seconds, the rest of the cast could improvise to fill the dead air.

But we can’t do that. It’s all pre-baked animation. Everything has to fit perfectly.

That’s very risky. And as crazy as we are, we knew we had to find ways to protect ourselves from human error when assembling everything.

Our design relies in encouraging player’s curiosity solely on the performance of actors.

So we did something rather unusual in games but it’s the bread & butter of theatre: **rehearsals.**

Rehearsals are extremely rare in games.

But they were crucial for us: great acting was one of our core pillars – and those performances had to be delivered within incredibly **strict timeframes.**

So, we spent a month rehearsing against a stopwatch; drilling scenes with the actors and tweaking the script until they fit within a margin of four seconds.

Rehearsals helped the actors to absorb not only their roles but the world they had to inhabitant, and understand the others’ motivations, conflicts and points of view.

It also allowed us to refine and modify the scripts or the layout based on ideas that emerged during rehearsal.

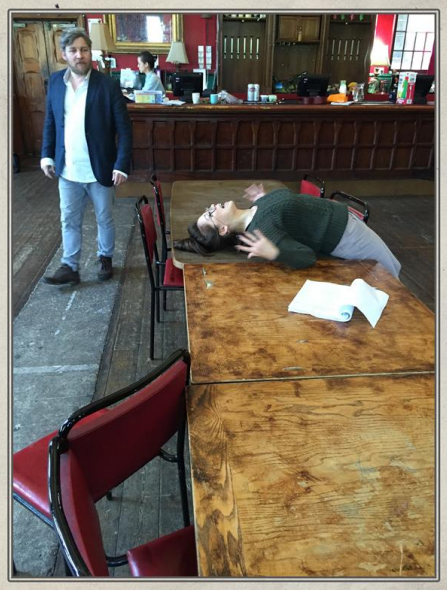

There is a risk when rehearsing this intensely, that the *humanity* is driven out of the performance
– so Rob spent a lot of time doing deep character exercises with the cast to help maintain the balance between efficiency and emotion.

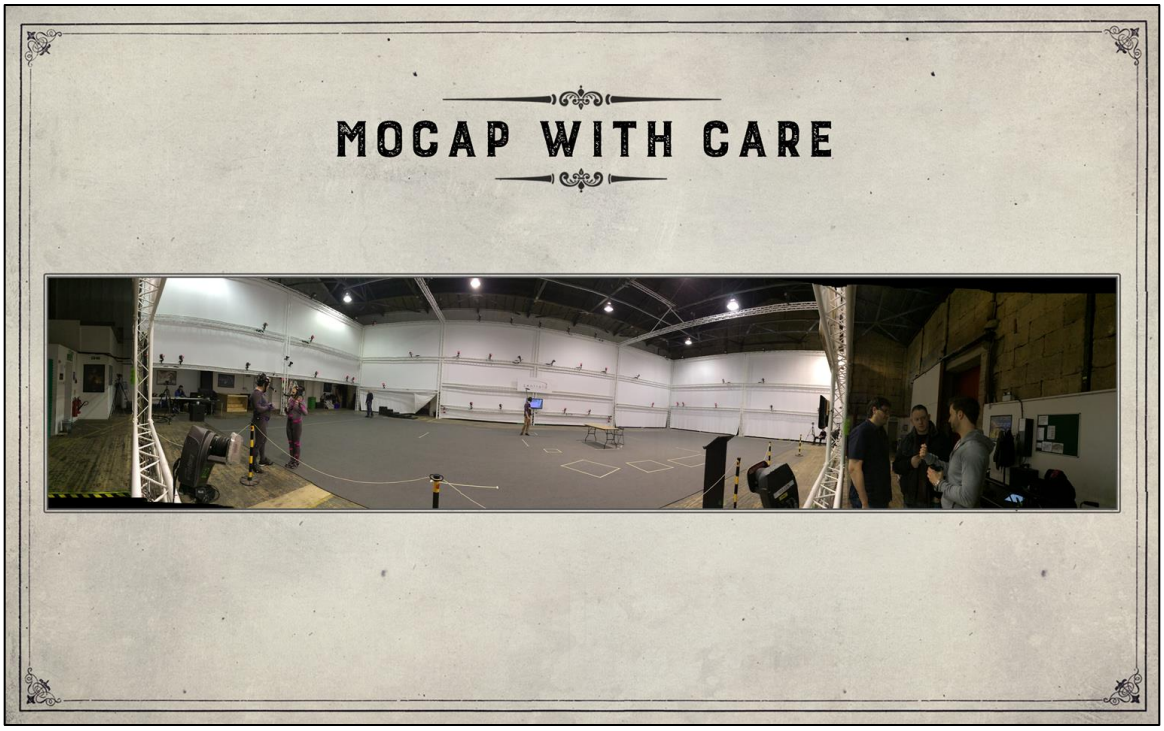

In a Spherical Narrative, the longer the story, the larger the margin for error becomes – so we split the story into four chapters.

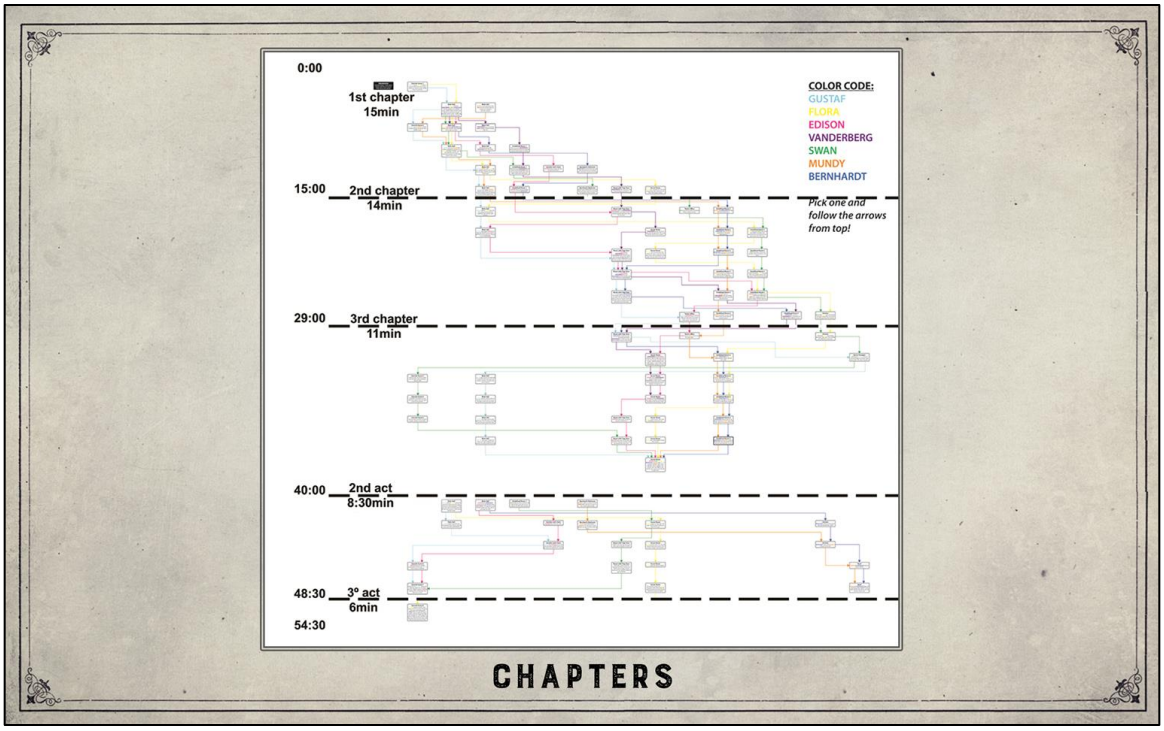

These chapters can be played back-to-back, but the separation between them gave us the opportunity to safely **resynchronize** everything if we discovered any unexpected time discrepancies.

These chapters also provide **natural** ‘intermissions’, just like in real theatre, where the audience can take a break for a while.

Motion and Facial Capture offer a balance between control and flexibility.

This left the door open to motion editing – so if a take happened to slip by a few seconds, we could subtly adjust the performance to bridge the gap.

Even with that additional flexibility, we had to manage and track our animation data very carefully: our four chapters are HUGE chunks of animation.

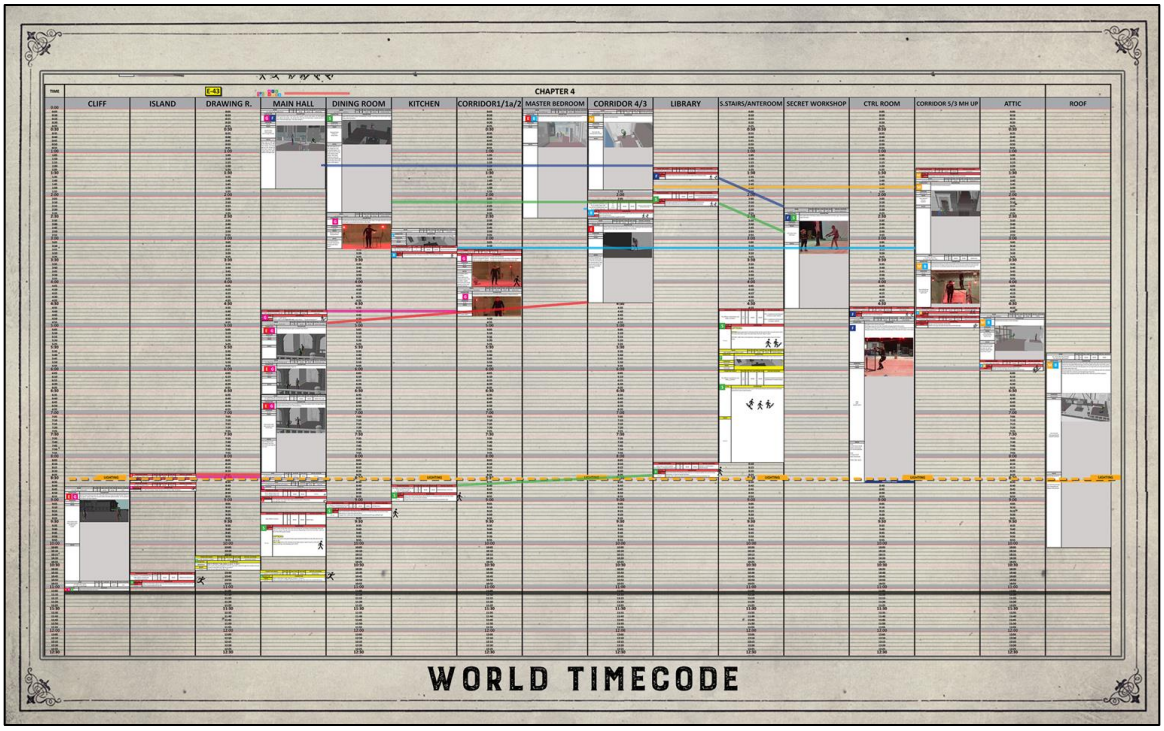

We're talking about almost 20 minutes of character data in Motion Builder *per chapter!* Each chapter alone is 80,000 frames multiplied per character.

Part of the challenge here was figuring out how to synchronize multiple facial and body data sets in co-existing scenes.
What we eventually realized was that we needed a **WORLD** timecode from the very first frame as zero, not a shoot timecode – which is what you'd usually do.

Like movies there were many takes taking into account different possible variations based on technical uncertainties or potential production bottlenecks.

If someone left a glass on a table that glass had to still be there when we re-assembled that set to shoot the next scene there three weeks later. Remember, there's no such thing as *off-screen*; you can't fix it later.

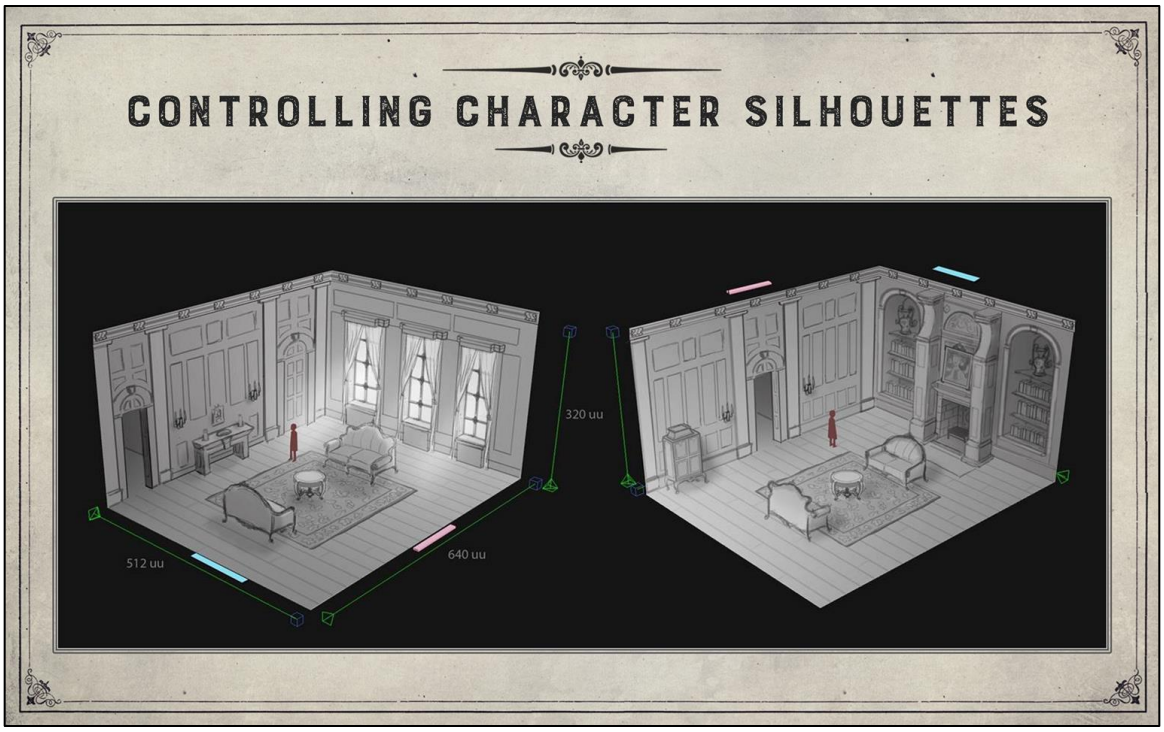

In normal animation, you make the character "perform for the camera" by posing them in the best way for that viewpoint. You can control character silhouettes.

But because we have no control over where or how people see animations, we had to direct a play that could be viewed from any angle.

It’s much easier to manage a character’s performance on a shot-by-shot basis – but here, characters are not just subjects of individual shots; they must be true and complete characters at all times.

Unlike “flat screen” experiences actors couldn’t know where the camera would be placed or the main spectator’s point of view as the player can be literally anywhere anytime.

This required the actors to rely more on improvisation and avoiding getting close to walls or anything that could block the way into rooms unless the script said so.

Each room became a stage where the actors would fill the empty spaces to perform so the player could move freely and observe at will.

Meanwhile, the team kept refining the whitebox, playing in VR the animatic with voiceover from the rehearsals.

Art defined the metrics and modular assets they would use to build the world, design experimented with the best way to create locomotion.

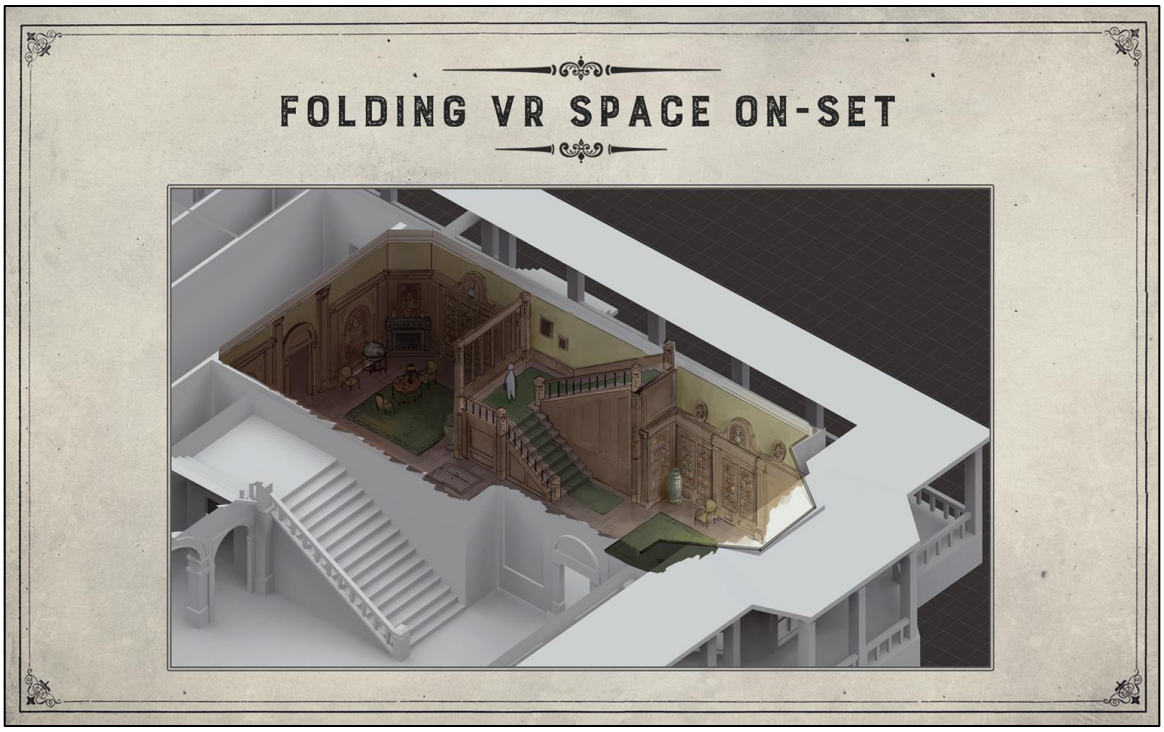

There were virtual rooms far bigger than our motion capture volume.

So, if you wanted two characters at opposite ends of the virtual room, they couldn’t both fit on the motion capture stage.

Shooting them separately risked both acting performance AND time syncing.

So instead, we did this:

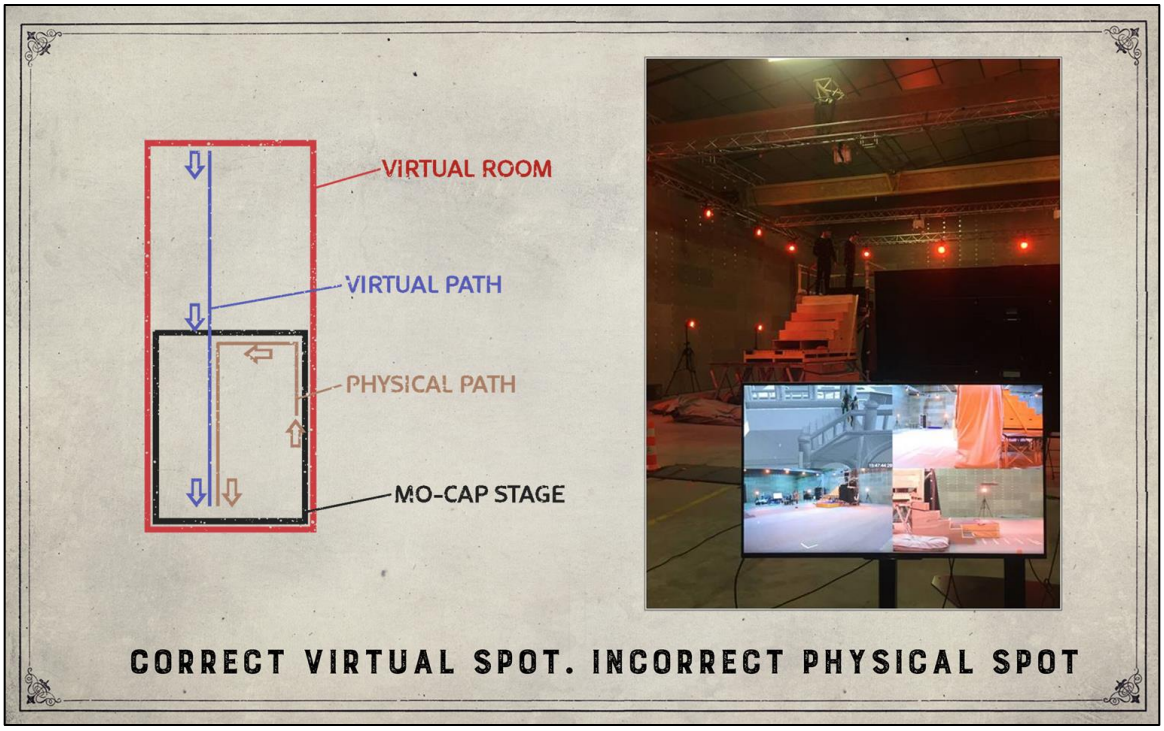

Folding the VR Space meant that our actors often had to perform to the correct virtual spot, but the incorrect physical spot in order for the pieces to fit together when we unfolded the space in-engine.

In the larger group scenes, this meant:

- keeping track of several different virtual **eyelines**
- and making sure that any movements in folded space kept in time with the rest of the scene

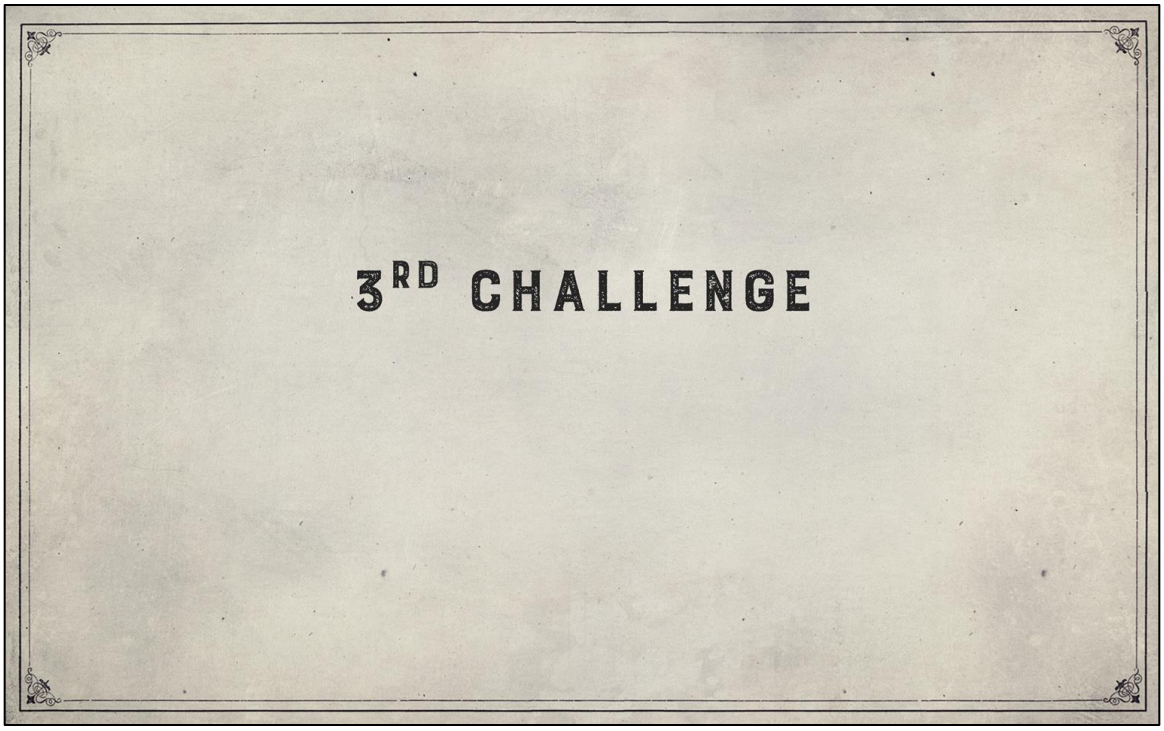

3rd Challenge

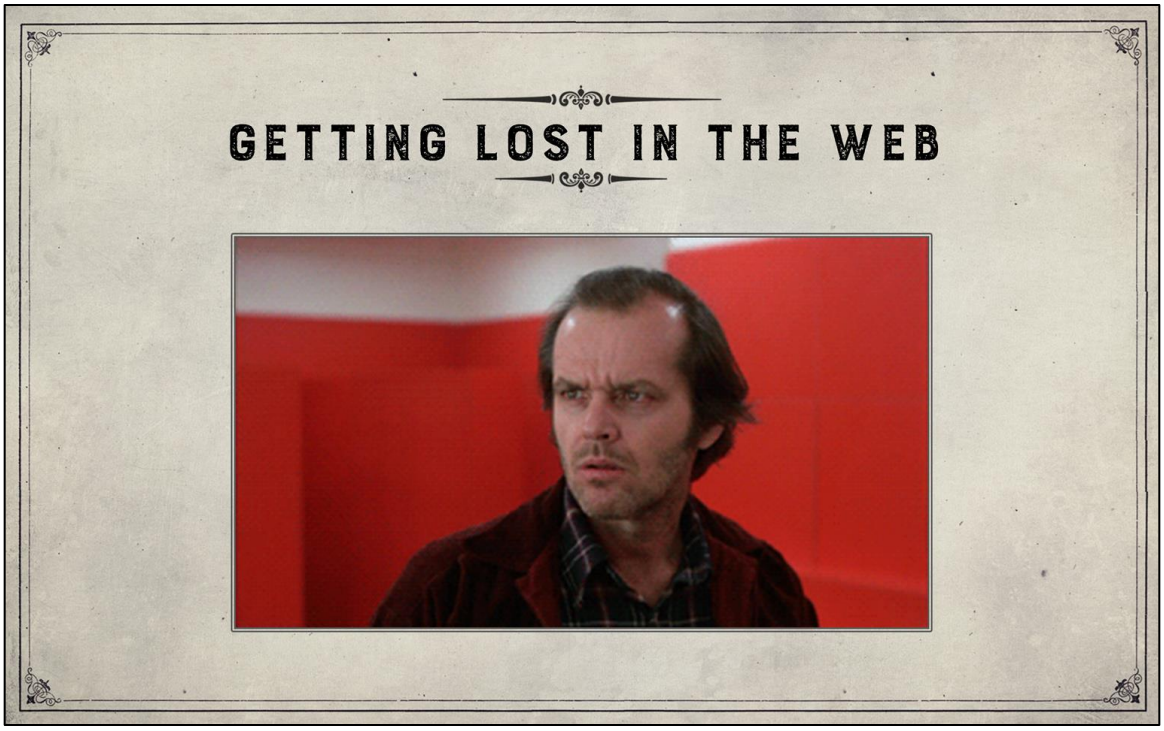

In most game narratives, the world is waiting for you to show up.

But in Spherical Narrative, it’s up to you to keep up with the world – just like real life.

But in doing that, there is a risk that the audience could lose the thread of the story they want to follow.

It’s *our* job to help them find it.

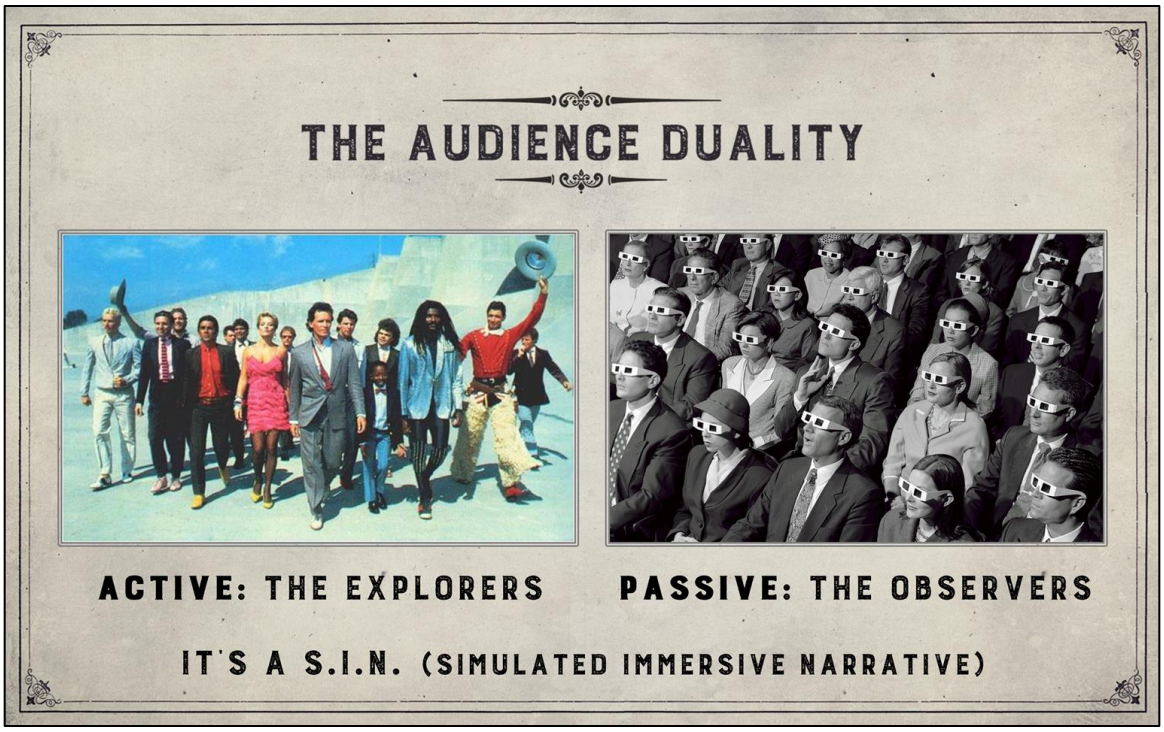

In addition to that, we have an audience duality.
This is not a game. This is not a movie. **It is a S.I.N. (Simulated Immersive Narrative).**
It had to appeal active audience or "explorers" (gamers) and passive **audience** or "observers" **(viewers**).
Our story had to work with gamers who feel comfortable with touch controllers, teleportation and control over time and space…
But also it couldn't stop Grandma to enjoy the experience like an episode of Angela Lansbury's *A Murder She Wrote*.
That means easy to navigate, easy to read, non-motion sickness inducing environments!
That means compelling, entertaining experience!
It had to be compelling **for Explorers,** the early adopters and har**d**core gamers who inve**sted on the device**.
It had to be accessible enough to be enjoyed by **our Grandma**, an Observer who won't even tu**rn his neck when wearing a HMD.**

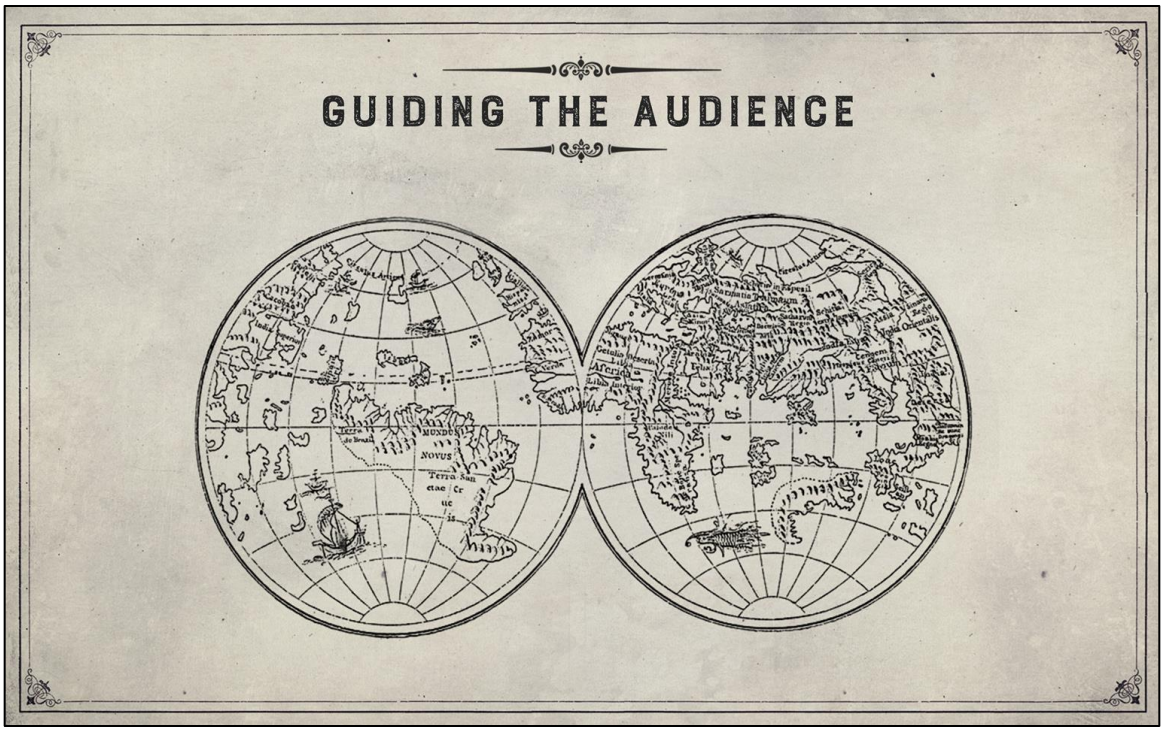

Unlike linear experiences, a spherical one is a full world in motion beyond your cognitive limitations.

Right now in this building alone there are countless stories that we won’t ever know in the story that it’s our live.

(For more detailed information see “*BUILDING THE STAGE FOR AN IMMERSIVE THEATER EXPERIENCE*“GDC talk by TW’s Artist Jaime González!)

Oscar’s Art Direction focused on removing all unnecessary visual noise to the experience, *amplifying* anything important to the story:
For example - any clues hidden should be easy to see, what can be interacted with.

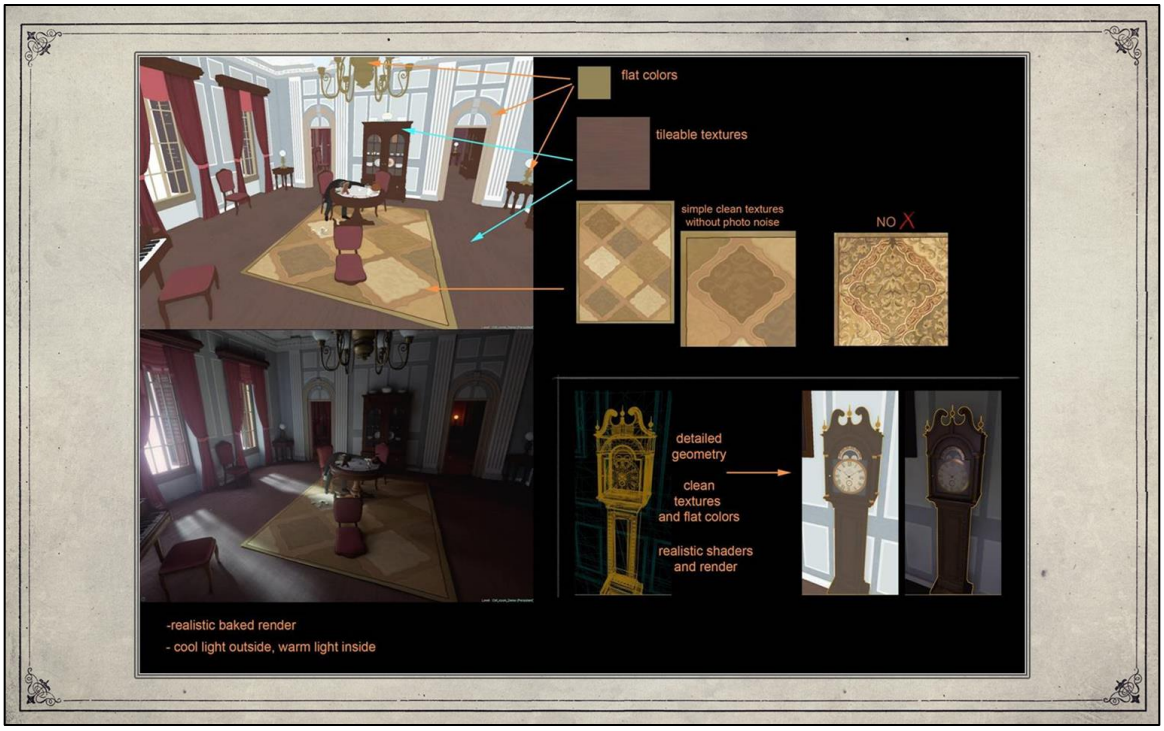

All is based on light and color, with stylized textures and characters that enhance shapes, facial animation and body language.

This is our way to guide the player and feed their curiosity

Also, having clear and distinctive character silhouettes made it easier to find the person you're looking for.

And in addition to that, by stylizing facial features, it became easier for players to see and absorb the characters' hidden emotions from a greater distance.

Audio is critical in VR. People working on VR already know that sound is a powerful tool for directing the player’s attention.

Audio is critical to let the player perceive events and situations beyond his field of view.

We encourage the player’s curiosity by the use of audio: in this project, lots of important events happen in locations beyond your line of sight.

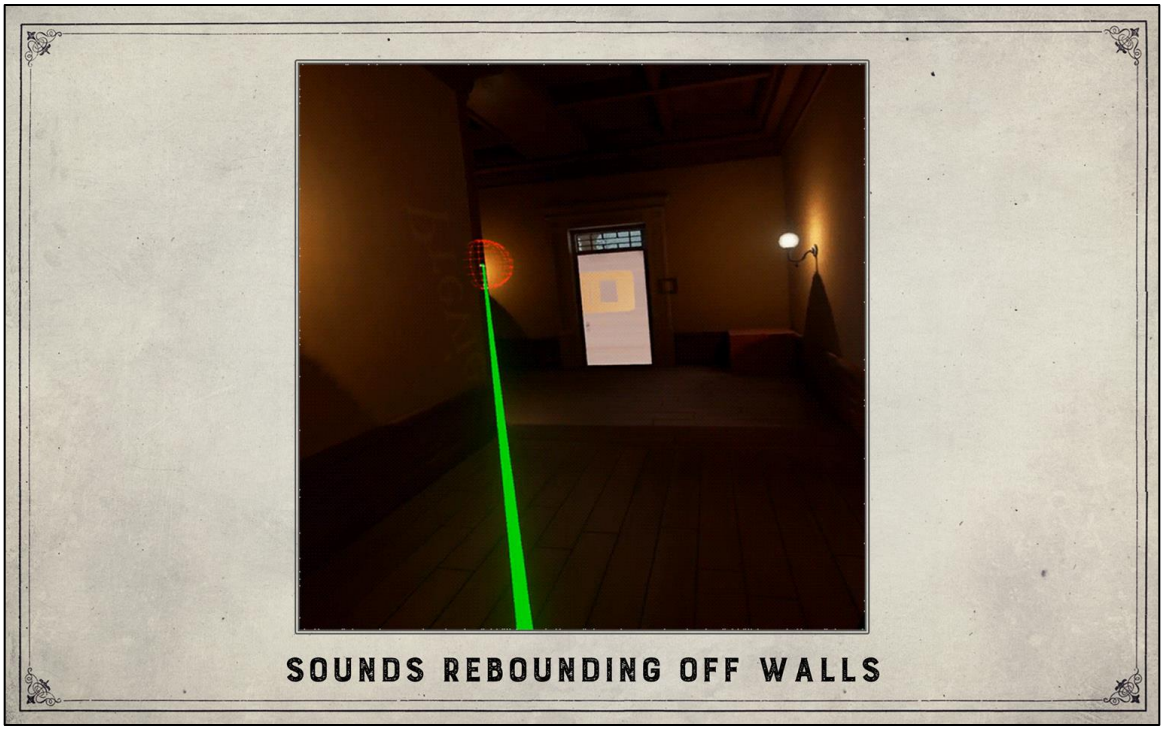

So, we created an audio system where sounds rebound off walls and hard surfaces.
So if there's a gunshot, or thunder, or if someone screams "murder!" the player hears it and knows what direction it came from and how far it happened – even if they're exploring some dark hidden corner of the mansion.

The music system was designed to balance the need to provide an emotional guide to the player VS making the player too aware of the artificiality of music in the world.
We composed the soundtrack not like a game or a movie, but a play. That means we use atmosphere all the time, not literal music tracks. But if the player approaches an important scene, the mood shifts and the music becomes a little more obvious.
This helps guide the player towards interesting moments around them.

Level Design is not only a way to create compelling gameplay.

It is the way to guide players so they don’t feel lost.
It is a way to reinforce their ingenuity and encourage the player’s curiosity to explore.

The world is telling us a story too. What happened is as important as what is happening.

The set provides information, clues and hints of past events.

If you do happen to find yourself standing in an empty room where nothing seems to be happening, there’s a layer of *environmental* narrative:
letters, diaries, and various clues are scattered around the house.

Maybe when the room is empty a secret passage opens.

Maybe a careful spectator will notice a subtle noise that could not be heard due to the conversation.

Forcing players to remember every location, character and event can be challenging.

Designers had to figure out how would players manage all this information.

The answer was the *Programme*: a 4D tracker of all events you have witnessed in time and space.

Think of it like the Marauder's Map from Harry Potter – but it only shows you what you've already seen.

So - it shows you all the paths of characters that you've followed, right up to the moment when you stopped following them.

That makes it easy to pick up a story thread again, by just returning to the time and place when you lost that particular character.

What's more, it also logs details you discover about the characters' relationships to each other and helps you understand their true motives and personalities.

The *Programme* is also an automatic way to track everything that the player interacts with.
This way, players can discover new things by connecting the dots by reviewing the action in 4D.

And of course, it's also a regular map – so players don't need to memorize the layout.

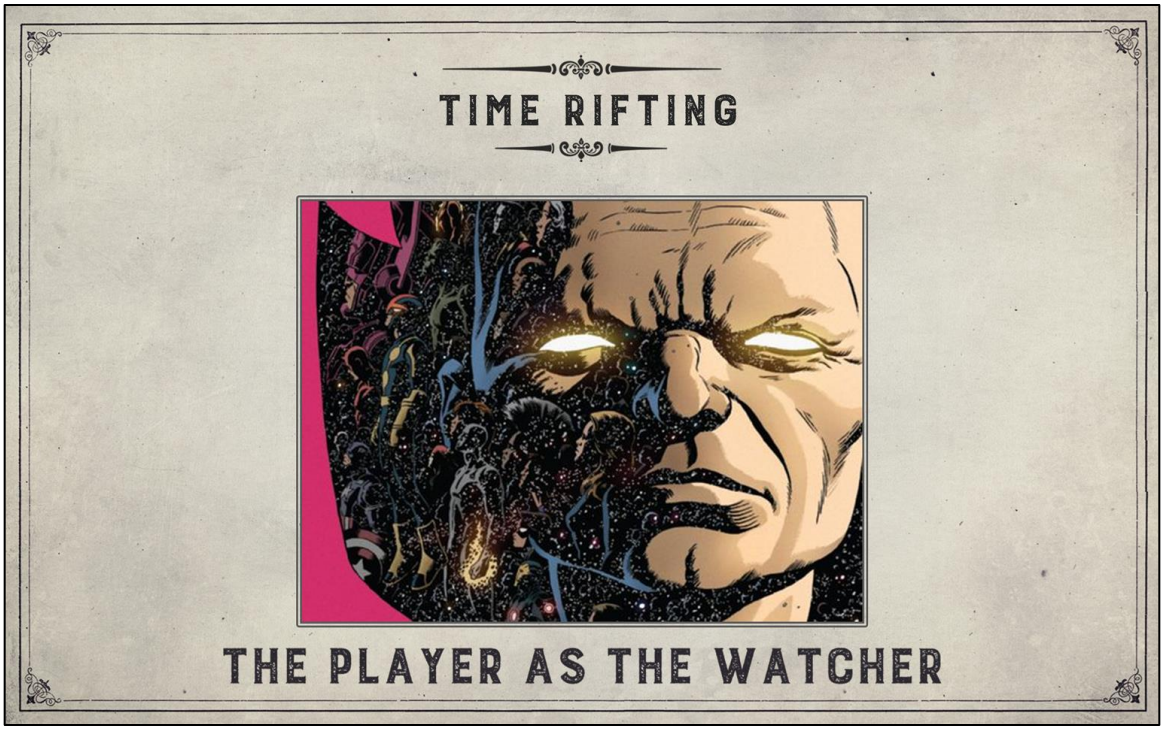

Unlike the characters in the story, the player is not restricted by time.
Time Rifting is a powerful 4D tool that also allows the player be the Director and play the part of the detective.
The more power we give players, the more likely they are to create a unique experience.
The player can revisit any location at any time to slowly complete the jigsaw.
In essence we want the player to be the Watcher of this story, the only being with the ability to gain omniscience and omnipresence.
The players are compensating their cognitive limitation in 3 dimensions by controlling the 4th.
With the Time Rifting players can:

- jump around any character's timeline
- revisit any scene and explore it in a different way
- or just rewind a few seconds to catch a line of dialogue they missed.

Or let's say you find a document burning – you can rewind time to see what that document said, then maybe rewind some more to see who burnt it, why, then maybe forward time again to see if they lie about it later.

Doing things like this can organically create moments of 'dramatic irony' that you would've experienced differently if you had watched the events in a different order.

We expect a lot of our players to be very active in their pursuit of the story – but of course, not everyone can do that.

What about your Grandma? Maybe she loves old re-runs of *MURDER SHE WROTE*, but asking her to go from simple TV to something this complex would be too much.

But we want her to play too, so we needed a system that would hold her hand.

So – by simply looking at any character, you can decide to ‘stick’ to them.

And if you do that, you will be gently teleported along with them wherever they go.

So, you can just sit back and have a passive experience if you want.

It’s almost like watching a regular movie.

It’s a less fulfilling experience, of course, but by trying the Grandma Mode first, maybe Grandma will feel comfortable enough to step off the track and explore for herself.

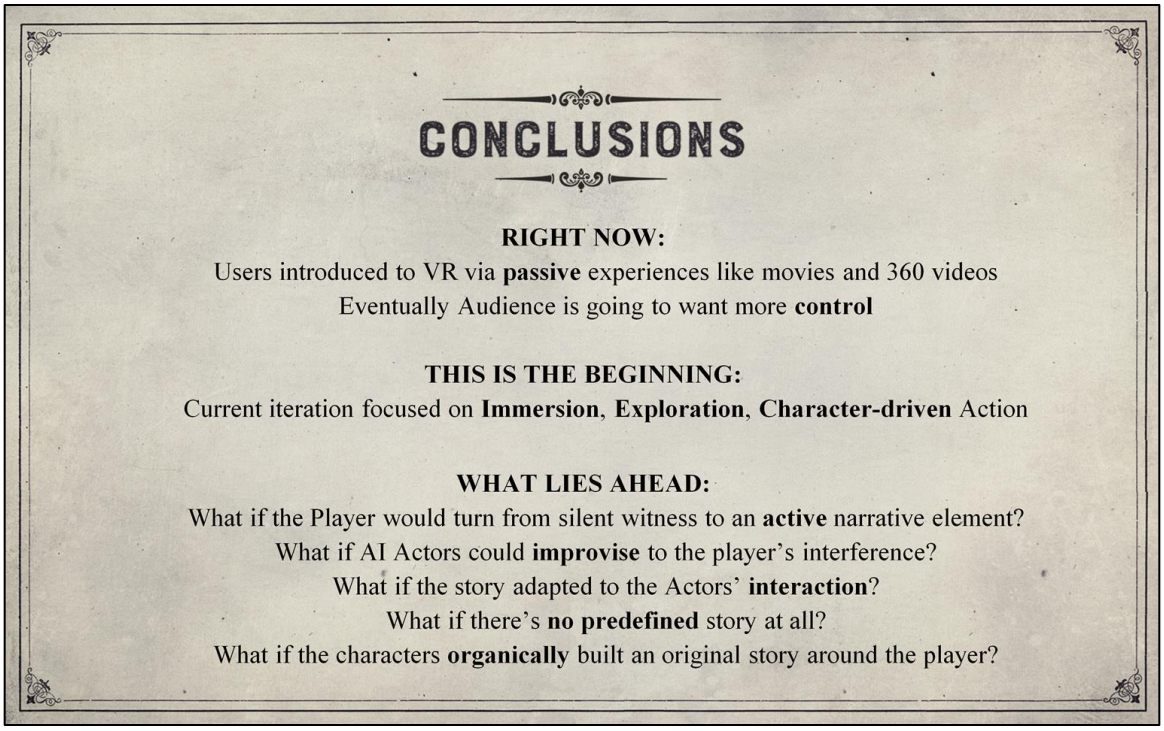

Right now, there are large numbers of VR users being introduced to the medium through very passive experiences – especially through things like VR movies and 360 videos.
But eventually a big part of that audience is going to want to take a little more control over their time inside VR stories – and Spherical Narrative is designed to give them that control.

This is just the beginning.
The current iteration of Spherical narrative focuses on immersion, exploration and character-driven action.
And, even though this particular project is a murder-mystery, we believe that the Spherical Narrative format has far wider-reaching potential.

What if we could make our AI to improvise?
What if they could really adapt to the player's presence and actions?
What if our characters would be really good performance artists and the story organically built around you?

– So please come and find us after the session, we’d love to talk more about it.

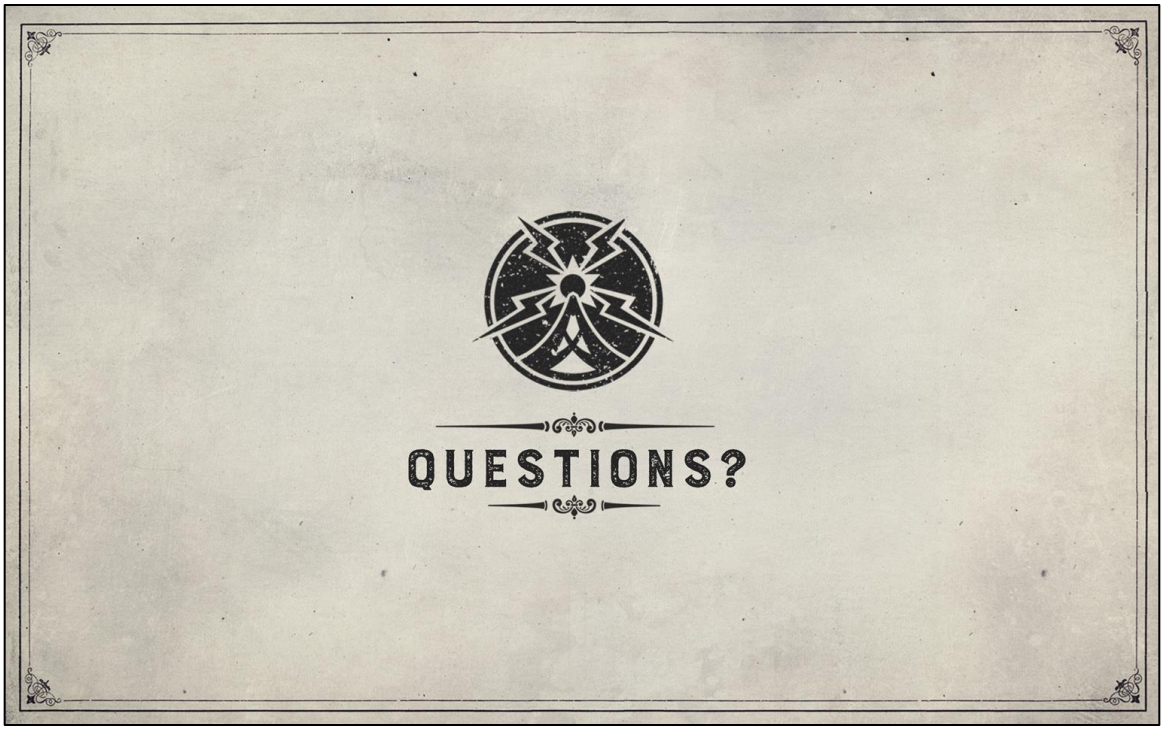

Questions?

TEQUILA WORKS