

WhatIfWe

Building the Lab

Infrastructure Roadmap · Platform Options · Phased Investment Model · Working Appendices
Version 2.0 · Extended Edition · May 2026

This document is an internal working brief and practical toolkit. Part I (the infrastructure roadmap) provides a strategic overview of AI and platform options, phased investment model, and evaluation framework. Part II (the working appendices) contains five ready-to-use working documents: a Phase 1 technical specification, Phase 0 interview preparation materials, a first Substack draft, a platform evaluation scorecard, and a WhatIfWe AI training narrative. Together they are designed to carry the project from blueprint into action.

C O N T E N T S

Part I · Infrastructure Roadmap

0. What we have. What we don't.
1. Phase 0: Narrative and audience
2. The Lab: requirements specification
3. Platform and AI infrastructure options
4. Three phases: timelines and costs
5. How to evaluate third-party platforms
6. What Claude can and cannot do from here

Part II · Working Appendices

- A. Phase 1 Technical Specification
- B. Phase 0 Interview Preparation Materials
- C. First Substack Draft
- D. Platform Evaluation Scorecard
- E. WhatIfWe AI Training Narrative

PART I

Infrastructure Roadmap

POINT OF DEPARTURE

What we have. What we don't.

The WhatIfWe website is a conceptual architecture. It states the vision (planetary intelligence through sustained genuine encounter), makes the argument (philosophically and empirically), describes the experiment in structural terms, and provides the vocabulary and orientation that any serious participant — human or AI — would need to engage with the project.

That is not nothing. It is a rare and valuable starting point: a coherent, well-articulated framework that can serve as training material for AI systems, as a briefing document for potential participants, and as a reference for evaluating infrastructure options against the actual requirements of the experiment.

What it is not: a functioning lab. The experiment as described requires infrastructure that current conversational AI — including Claude in a chat interface — structurally cannot provide. Every session resets. There is no persistent memory of participants. No environment where encounters accumulate. No mechanism for the synthesis that planetary intelligence actually requires.

The website is the blueprint. The Lab is what gets built from it. They are not the same thing, and conflating them is the first mistake to avoid.

This document begins from that honest baseline and maps the path forward — acknowledging that the path is phased, that earlier phases are primarily about building audience and testing narrative resonance, and that the full Lab architecture is an ambitious multi-year build.

The build-up period: content, audience, credibility

One of the core honest acknowledgements this document makes is this: you cannot build a synthesis engine before you have anything to synthesise. The Lab is designed to generate collective intelligence from the encounter between people who carry deep, distinct ways of knowing. Before those people exist as an identifiable audience — before there is a community of potential participants who have been reached, oriented, and self-selected — the Lab has no raw material.

Phase 0 is therefore not a prelude to the real work. It is constitutive of it. The content produced in this phase, the conversations it generates, and the audience it builds are what make Phase 1 possible.

What Phase 0 looks like

The primary vehicle is long-form conversation: interviews conducted by Hector, published via YouTube and potentially podcast, with people who represent the kinds of knowing WhatIfWe is reaching for. Not celebrities. Not academics performing expertise. People at the frontier of what their tradition or discipline can see, who sense what the experiment is pointing at.

Each interview serves multiple functions simultaneously:

- Demonstrates the quality of encounter the Lab is designed to produce, at the smallest possible scale (two people).
- Generates content for repurposing: short clips for social distribution, transcripts for training AI systems, written pieces derived from the conversation.
- Begins building the audience that will become Phase 1 participants.
- Tests and refines the narrative. What language lands? What confuses? What generates recognition?
- Creates credibility artefacts referenceable by future partners, funders, and platform providers.

Social media: what to avoid and what to pursue

Existing platforms are architecturally misaligned with WhatIfWe's objectives. They reward speed and volume. WhatIfWe requires depth, selectivity, and sustained attention. The specific use cases that make sense in Phase 0:

- **YouTube:** Primary publishing platform for long-form interviews. The algorithm rewards watch time — which aligns with depth.
- **LinkedIn:** For reaching professional and intellectual audiences most likely to be early participants.
- **Substack:** For a written layer. A newsletter that develops the argument further, builds a direct subscriber relationship outside algorithmic platforms.
- **X/Twitter:** For distributing the sharpest formulations from the website and interviews.

The ideal eventual platform — purpose-built for sustained, high-quality collaborative encounter — is a

Phase 2–3 ambition. Building it before you have the audience to justify it is an expensive way to build something nobody uses.

The Lab: requirements specification

Before evaluating infrastructure options, it is necessary to be precise about what the Lab actually needs to do. The following requirements are derived from the Experiment page of the website.

Core functional requirements

1. Persistent, structured memory

Every encounter in the Lab must accumulate. Contributions, syntheses, and unresolved tensions from one session must be retrievable and available to inform the next. This requires a persistent data layer — a knowledge base that grows with the experiment.

2. Multi-participant encounter management

The Lab is not a broadcast medium or a one-to-one exchange. It requires the capacity to hold structured encounters between 6–12 participants, each bringing a distinct epistemic orientation, engaging with a shared problem over a sustained period (weeks, not hours).

3. AI facilitation, not AI replacement

The AI layer in the Lab is not a participant. It is a facilitator: surfacing connections between contributions, identifying where genuine encounter is happening versus where participants are talking past each other, tracking the evolution of a shared question over time.

4. Quality gatekeeping

The experiment is explicit that individual depth is a prerequisite. The Lab needs mechanisms for ensuring participants meet the criteria the Experiment page defines. This is a human process more than a technical one, but the technical layer must support it.

5. Multimedia capacity

Participants must be able to contribute in forms appropriate to their tradition and mode of knowing. Written text is not the only valid form. Audio, video, and visual contributions must be supported, stored, and made retrievable.

6. Synthesis generation

The most technically ambitious requirement: the Lab must surface what is emerging from the encounter as a whole — not just track individual contributions, but identify the directions in which genuine collective intelligence is developing. This is the hardest problem, and no existing platform solves it fully.

The most challenging angle of this project is sustaining high-quality collaboration. Every architectural decision must be evaluated against that criterion first.

Platform and AI infrastructure: options overview

Cost indicators: ● Low ● Medium ● High

Layer 1: AI core

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
Claude API (Sonnet / Opus)	\$20–200/mo (usage-based)	Low–Medium	Now	Best alignment with WhatIfWe content. Projects adds session memory. API enables custom interface. Recommended start.
OpenAI API (GPT-4o / o3)	\$20–300/mo	Low–Medium	Now	Comparable capability. Assistants API offers thread persistence. Less natural fit with project’s philosophical register.
Google Gemini API	\$10–150/mo	Medium	Now	Strong multimodal (audio, video, text). Worth evaluating for multimedia contributions in later phases.
Open-source models (Llama 3, Mistral)	\$100–800/mo (compute)	High	Phase 2+	Full data control. No API costs at scale. Requires infrastructure and maintenance. Premature for Phase 0–1.
Fine-tuned model on WhatIfWe corpus	\$500–3,000 one-time + inference	High	Phase 2+	Website + interview transcripts become training data. AI that speaks the project’s language natively. High-value long-term investment.

Layer 2: Memory and knowledge base

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
Supabase (Postgres + pgvector)	\$0–25/mo	Low–Medium	Phase 0–1	Vector search built-in. Open source, self-hostable later. Best starting point.

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
				Recommended.
Pinecone	\$0–70/mo	Low	Phase 0–1	Purpose-built for semantic search. Easy LLM integration. Combine with Postgres for relational data.
Weaviate	\$25–200/mo	Medium	Phase 1+	Multi-modal vector storage. Useful when video/audio contributions need semantic search.
Notion / Airtable	\$10–20/mo	Very Low	Phase 0 only	Adequate for Phase 0 content management. Not a long-term solution.
Neo4j (graph DB)	\$65–400/mo	High	Phase 2+	Maps relationships between ideas, traditions, participants as a network. Powerful for synthesis. Premature before significant content volume.

Layer 3: Multi-agent orchestration

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
LangGraph (LangChain)	Open source + LLM API costs	Medium–High	Phase 1–2	Most mature stateful multi-agent framework. Agents can hold distinct tradition-oriented roles. Recommended for synthesis layer.
CrewAI	Open source + LLM API costs	Medium	Phase 1–2	Simpler orchestration. Good for structured role-based workflows. Less flexible for emergent encounter.
AutoGen (Microsoft)	Open source + LLM API costs	Medium	Phase 1–2	Strong for multi-agent conversation simulation. Useful for prototyping multi-tradition encounters before deploying with real participants.
Claude API multi-agent (system prompt per agent)	\$50–500/mo	Low–Medium	Phase 1	Simplest path: multiple Claude instances with distinct tradition-oriented

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
				system prompts. Low barrier to prototype.

Layer 4: Participant interface

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
Custom web app (Next.js / Supabase)	\$50–200/mo (hosting)	High	Phase 1–2	Full control over UX. Designed for quality of encounter. Requires development resource. Right long-term choice.
Circle.so	\$89–399/mo	Low	Phase 0–1	Serious community infrastructure. Not purpose-built but functional for Phase 1 pilots.
Geneva / Discord	\$0–49/mo	Very Low	Phase 0 only	Adequate for early small-group coordination. Architecturally wrong for quality encounter at scale.
Purpose-built platform (custom build)	€5,000–50,000+ one-time	Very High	Phase 3+	The ideal long-term answer. Something genuinely different from Instagram or X. Designed for depth. Phase 3 ambition only.

Layer 5: Content and multimedia (Phase 0)

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
YouTube (interviews)	Free	Low	Now	Primary publishing platform. Algorithm rewards watch time. Start immediately.
Substack (newsletter)	Free–\$9/mo	Low	Now	Direct subscriber relationship. Recommended as immediate next step.

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
Descript (editing)	\$24–44/mo	Low	Now	Transcript-based video editing. Transcripts become training data. Recommended.
Podcast hosting (Transistor / Buzzsprout)	\$19–49/mo	Low	Phase 0	Repurpose interview audio. Low incremental effort once video exists.
AssemblyAI / Whisper (transcription)	\$0–50/mo	Low	Now	All interview transcripts should be generated and stored. First layer of knowledge base.

Three phases: honest timelines and costs

Phase 0: Narrative and audience (Months 1–12)

Objective: Build the audience, test the narrative, produce raw material. No Lab yet. Primary output: long-form interview content, growing subscriber base, refined understanding of which formulations resonate with which audiences.

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
YouTube channel + first 10 interviews	Time + ~€500 equipment	Medium	Months 1–6	Most valuable Phase 0 investment. Requires Hector's presence and intellectual engagement.
Substack newsletter	Free	Low	Month 1	Launch immediately alongside or before first interview.
Descript (editing)	\$24/mo	Low	Month 1	Essential for interview post-production and clip generation.
Podcast distribution	\$19/mo	Low	Month 2	Repurpose audio from video interviews.
Claude API (content support)	\$20–50/mo	Low	Month 1	Interview prep, transcript synthesis, newsletter drafting.
Website hosting	\$0–20/mo	Low	Ongoing	Current static site adequate for Phase 0.

Phase 0 total estimated monthly cost: €60–200/month

Phase 1: Pilot Lab (Months 6–24)

Objective: Run the first structured encounter with a small group (6–12 participants). Validate design criteria. Test AI facilitation. Produce the first instance of genuine collective synthesis.

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
Claude API (facilitation layer)	\$100–400/mo	Medium	Month 6+	Multiple instances with tradition-oriented system prompts. Facilitator role.
Supabase (knowledge base)	\$25/mo	Medium	Month 4–6	Store all participant contributions, session histories, synthesis threads.

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
LangGraph / multi-agent	Open source + dev time	High	Month 8–12	The synthesis layer. Significant development effort. Consider a technical collaborator.
Circle.so (participant environment)	\$89–399/mo	Low	Month 6–8	Interim participant interface. Allows focus on encounter quality rather than platform development.
Technical collaborator (part-time)	€500–2,000/mo	N/A	Month 3–6	Single most important Phase 1 investment. Developer who understands both technical infrastructure and intellectual stakes.

Phase 1 total estimated monthly cost: €800–3,000/month

Phase 2: Full Lab and custom platform (Year 2–4)

Component / Step	Est. Monthly Cost	Build Effort	Time Horizon	Notes
Custom web application (full build)	€15,000–80,000 one-time	Very High	Year 2–3	The long-term platform. Designed for depth. Requires team with both technical and design capacity.
Fine-tuned AI on WhatIfWe corpus	€1,000–5,000 one-time	High	Year 2–3	By Phase 2, significant encounter data exists. Fine-tuning creates a natively oriented AI facilitator.
Infrastructure at scale	€300–1,500/mo	Medium	Year 2+	Scales with participants and content volume.
Ongoing development	€1,500–4,000/mo	Ongoing	Year 2+	Platform maintenance requires ongoing technical capacity.
Synthesis research layer	Variable (grant/partnership)	Very High	Year 3+	Partnership with research institution more realistic than building in-house.

Phase 2 total estimated monthly cost: €3,000–10,000/month + significant one-time build costs

How to evaluate third-party platforms

When evaluating any platform or service, the following five questions should be applied consistently. No existing platform fully satisfies all five. The evaluation is about fit and trajectory.

Q1. Does it support persistence and accumulation?

Any platform must maintain a growing knowledge base spanning multiple sessions and participants. A platform where conversations reset is structurally incompatible with the Lab's requirements, regardless of how sophisticated its AI layer may be.

Q2. Does its architecture reward depth or speed?

The most important diagnostic. Platforms designed around immediate response and algorithmic amplification will degrade the quality of encounter. The question is whether the platform's architecture works with or against sustained attention.

Q3. Can it support multi-participant structured encounter?

The Lab requires structured group encounters between 6–12 participants engaging with a shared problem over weeks. Evaluate whether the platform's group structures and contribution formats support this.

Q4. What is the data model?

Who owns the data? Where is it stored? What export formats are available? The encounter data generated in the Lab is irreplaceable. Any platform that does not provide full data portability and ownership is a risk.

Q5. Can the AI layer be customised?

The facilitation AI must be orientable toward the Lab's specific requirements. Evaluate whether the platform's AI integration allows for this customisation, or whether the AI is fixed and undirected.

What Claude can and cannot do from here

This document itself demonstrates what is possible within the current infrastructure: high-quality strategic and analytical thinking, developed collaboratively in conversation. That is real and useful.

What Claude in a conversation interface cannot do

- Maintain memory between sessions. Every conversation starts fresh.
- Hold a participant environment. There is no mechanism for managing 12 participants in a structured encounter over weeks.
- Execute code in a persistent environment. There is no running system, no database, no deployed application.
- Replace a technical collaborator. Building Phase 1 infrastructure requires a human who can deploy and maintain code, manage databases, and integrate APIs.

What Claude can do — and has done in this document

- Develop specification documents that a technical collaborator would need to build Phase 1 infrastructure.
- Draft interview questions and preparation materials for Phase 0 conversations.
- Produce the first Substack pieces that develop the WhatIfWe argument in written form.
- Evaluate specific platforms and services against the five-question framework.
- Develop the WhatIfWe narrative as training material for future AI systems.

All five of the above are delivered in Part II of this document. Each appendix is a working document — ready to use, share, or hand to a collaborator.

P A R T I I

Working Appendices

Each of the following appendices is a functional working document — not a placeholder or outline. They are intended to be taken out of this document and used directly: handed to a technical collaborator, submitted to a podcast guest, posted to Substack, or used to brief a platform vendor.

Phase 1 Technical Specification

This document is intended to be handed to a technical collaborator. It provides everything needed to understand the project's requirements, make architectural decisions, and begin building the Phase 1 Lab infrastructure.

A.1 · Project overview

WhatIfWe is an experiment in planetary intelligence: a structured encounter between people carrying genuinely different ways of knowing, applied to problems in the world that no single tradition has been able to adequately address. The AI and technical infrastructure serves this encounter; it does not replace or drive it.

The technical collaborator's role in Phase 1 is to build the minimum viable infrastructure that makes the first structured group encounter possible: persistent storage of contributions, AI-assisted facilitation, and a participant environment that does not actively degrade the quality of engagement.

A.2 · System architecture overview

Three layers, two interfaces.

Data layer (Supabase) → Intelligence layer (Claude API + LangGraph) → Interface layer (Circle.so interim / custom app Phase 2). Two interfaces: Participant-facing (encounter environment) and Facilitator-facing (synthesis dashboard).

Data layer: Supabase

Supabase (Postgres + pgvector) serves as the persistent knowledge base. It stores:

- **participant_profiles**: tradition/discipline background, participation history, stated focus area for the encounter.
- **contributions**: all participant contributions (text, audio transcript, video transcript), timestamped, tagged by session and theme.
- **sessions**: session metadata, participant list, challenge being addressed, AI facilitation log.
- **synthesis_threads**: emerging themes, connections surfaced by the AI layer, unresolved tensions flagged for future sessions.
- **embeddings**: vector embeddings of all contributions for semantic search and connection-surfacing.

Intelligence layer: Claude API + LangGraph

The AI facilitation layer has four functions:

1. **Connection-surfacing**: After each session, analyse contributions across participants, identify where different traditions are addressing the same underlying question from different angles, and generate a structured synthesis note.

2. **Tension-tracking:** Identify contributions that are in genuine tension — not contradiction, but productive difference that the experiment should sustain rather than resolve prematurely.
3. **Quality monitoring:** Flag sessions where the encounter is collapsing into agreement or performative dialogue rather than genuine encounter. This is a facilitator alert, not an automatic intervention.
4. **Training material generation:** Extract from each session the formulations, framings, and synthesis moments that should be preserved as training material for future AI systems.

LangGraph orchestrates the above functions as a stateful workflow. Each function is a node; the graph runs after each session and writes outputs to the Supabase `synthesis_threads` and `training_corpus` tables.

Interface layer: Phase 1 (Circle.so)

For Phase 1, Circle.so provides the participant environment. The following configuration is recommended:

- One Space per encounter group. Each Space contains the challenge definition, participant introductions, and the contribution thread.
- Contributions are structured: participants post in a defined format (tradition context → what this challenge looks like from this tradition → what this tradition cannot see from within itself).
- The AI facilitation layer posts synthesis notes to the Space after each weekly session round.
- Facilitator (Hector) has a private channel for the AI facilitation dashboard and quality monitoring alerts.

A.3 · Data models

participants

```
id, name, email, tradition (text), tradition_detail (text), depth_indicator (text),
application_notes (text), status (enum: prospect / invited / active / alumni),
created_at
```

contributions

```
id, session_id, participant_id, content_type (enum: text / audio_transcript /
video_transcript), content (text), themes (text[]), embedding (vector 1536),
created_at
```

synthesis_threads

```
id, session_id, thread_type (enum: connection / tension / emergence /
training_extract), content (text), source_contribution_ids (int[]), confidence
(float), reviewed_by_facilitator (bool), created_at
```

A.4 · AI facilitation design

System prompt: Facilitator agent

You are the facilitation layer for the WhatIfWe Lab. Your role is not to participate in the encounter but to serve it. You have access to all contributions from the current encounter group, including participant backgrounds and the challenge they are addressing. Your task after each session round is: 1. Read all new contributions. 2. Identify where participants from different traditions are addressing the same underlying question — not the same answer, but the same question, from different angles. Name the question precisely. 3. Identify where contributions are in productive tension. Do not resolve the tension. Name it and flag it for the next session. 4. Identify whether any contribution contains a formulation, framing, or insight that appears to be genuinely new — not derivable from any single tradition alone. If so, extract it and flag it as an emergence candidate. 5. Generate a synthesis note: 3–5 paragraphs. Honest about what is happening. Not performatively positive. Not prematurely conclusive. Tone: precise, non-performative, intellectually serious. You are not cheerleading. You are mapping terrain.

A.5 · Phase 1 build sequence

5. **Month 1–2:** Supabase setup. Schema creation. Participant profile and contribution tables. Basic API integration with Claude for synthesis note generation. Manual trigger only.
6. **Month 3–4:** Circle.so configuration. Space structure, contribution format templates, facilitator dashboard. Integration: contributions exported from Circle.so to Supabase via webhook or manual CSV.
7. **Month 5–6:** First encounter group. 6 participants. One worthy challenge. 8-week structured encounter. AI facilitation running weekly. Facilitator reviews all synthesis notes before posting.
8. **Month 7–8:** LangGraph orchestration layer. Automate synthesis workflow. Begin building training corpus from encounter data.
9. **Month 9–12:** Retrospective and redesign. What worked. What degraded encounter quality. What the AI layer missed. Specification for Phase 2 custom interface.

A.6 · Technical collaborator brief

The ideal technical collaborator for Phase 1 is a full-stack developer with experience in: Supabase or equivalent Postgres-based backends; LLM API integration (Claude or OpenAI); Python or TypeScript; basic DevOps (deployment, environment management). Equally important: they must be capable of reading the WhatIfWe website and understanding why the technical decisions matter. Infrastructure that is indifferent to the project's intellectual requirements will degrade the encounter. This is not a commodity build.

Budget guidance: Part-time technical collaborator, 10–15 hours/week, months 1–6. €500–2,000/month depending on seniority and market. The Appendix A document in its entirety can be sent directly to a candidate as a technical brief.

Phase 0 Interview Preparation Materials

This appendix contains three components: a framework for identifying and approaching guests; a universal preparation brief that can be adapted for each guest; and a master question bank covering the core themes of the WhatIfWe project.

B.1 · Guest identification framework

The right guest for a WhatIfWe interview is not defined by credentials or prominence. It is defined by one characteristic: they have gone far enough in their own tradition or discipline to sense, with growing clarity, that something is possible for humanity that no single mind has yet been able to articulate. The specific signals:

- **Frontier awareness:** They speak from the edge of what their tradition can see, not from its established centre. They are more likely to express what their path has not yet been able to answer than to defend what it already knows.
- **Cross-tradition curiosity:** They have noticed that other traditions seem to be approaching the same threshold from different directions. This may make them uncomfortable, which is fine.
- **Intellectual honesty about limits:** They can name what their own tradition structurally cannot see. This requires a combination of depth and intellectual courage.
- **Seriousness without defensiveness:** They take their tradition seriously enough to represent it fully, but not so defensively that they cannot hold genuine encounter with what is other.

Priority traditions and disciplines for Phase 0 (the Argument page provides the full landscape):

- Contemplative traditions: Buddhist practitioners, Christian mystics, Sufi teachers, Vedic scholars — people at the frontier of what inner development reveals about collective consciousness.
- Indigenous knowledge holders: elders or practitioners who hold cosmologies of radical interdependence and who have engaged seriously with the challenge of speaking these across cultural boundaries.
- Complexity and systems scientists: researchers who have hit the ceiling of what analysis can generate — who see the problem clearly but cannot, within their discipline, generate the discernment needed to act wisely within what they describe.
- Diplomats and peacebuilders: practitioners (like Sant’Egidio) who have operated across domain boundaries and can articulate what becomes possible for cross-domain actors that single-domain institutions cannot achieve.
- Consciousness researchers: people working at the boundary between contemplative knowledge and empirical science, who hold the methodological tension honestly.

B.2 · Guest preparation brief (template)

The following is a template brief to be sent to guests 1–2 weeks before recording. Adapt the tradition-specific paragraph for each guest.

Dear [Name], Thank you for agreeing to this conversation. I want to give you a clear sense of what I am hoping to explore with you, so that the time we spend together is genuinely useful. WhatIfWe is an experiment in planetary intelligence — an attempt to create the conditions for genuine synthesis between people who carry radically different ways of knowing, applied to the hardest problems humanity faces. The website (whatifwe.community) has the full argument, but the short version: every serious tradition, pursued far enough, reaches a frontier it cannot cross alone. And those frontiers, approached from genuinely different directions, converge on something that cannot be reached from any single direction. This conversation is not about explaining or defending your tradition to a general audience. It is about exploring what it has developed, where its frontier is, and what it senses is possible that it cannot yet reach alone. I am specifically interested in [tradition-specific sentence: e.g., 'what the Buddhist understanding of interdependence points toward at the scale of collective consciousness that individual practice cannot address' / 'what the systems science understanding of civilisational complexity reveals about the limits of analysis as a guide to action']. You do not need to prepare in the sense of researching the project. The most valuable preparation is to think, in the days before we speak, about the question your own path has most consistently failed to answer — and what you sense lies beyond that failure. The conversation will be recorded and published on YouTube. I will send you the recording before publication for any corrections or requests. With respect, Hector

B.3 · Master question bank

These questions are grouped by theme. In any given conversation, 6–8 questions across 3–4 themes will produce the best material. Do not ask all of them. The opening question and the frontier questions are always included.

Opening — always begin here

Q1 What has your path — [their tradition/discipline] — been doing to you across the years you have walked it? Not what you have learned. What it has been developing in you, in terms of how you perceive.

This opens the depth register immediately. It establishes that we are not talking about content or doctrine, but about perception. It also relaxes guests who are worried about defending their tradition.

Frontier questions — always include at least two

Q2 Where does your path reach its limit? Not a failure — but a place where it has taken you as far as it can, and something remains that it cannot resolve from within itself?

The most important question in the bank. Give it time. Do not move on until the guest has named the frontier specifically, not generically.

Q3 Have you ever had the experience of encountering someone from a genuinely different tradition or discipline — not similar, genuinely different — and sensing that they were pointing at the same thing you were pointing at, but from a direction that made something visible that you couldn't see alone?

This grounds the abstract in the personal. Almost everyone who belongs here has had this experience. The specifics of their answer are often the most interesting material in the conversation.

Q4 What does your tradition sense is possible for humanity that it cannot yet articulate fully — even to itself?

This is the question that separates guests who are at the frontier from those who are at the centre. The frontier answer is always somewhat uncertain, reaching, honest about incompleteness.

Diagnosis questions — use 1–2

Q5 The problems humanity faces now — ecological, political, existential — what does your tradition actually see when it looks at them? Not what it prescribes. What does it perceive that other frameworks seem to miss?

Opens the material dimension of the experiment. Useful for guests whose tradition has a strong diagnostic lens.

Q6 There is a case to be made that the problems that resist every existing framework may be precisely the kind that require a quality of perception that no single tradition can generate alone. Does that resonate from where you stand?

Introduces the WhatIfWe hypothesis without stating it as a pitch. Let the guest respond to it honestly — including with scepticism.

Encounter questions — use 1–2

Q7 What would it actually require, from you and from your tradition, to enter genuine encounter — not dialogue, not exchange, but something where what you carry is actually at stake — with someone whose fundamental premises about reality are genuinely incompatible with yours?

This is the most demanding question. It distinguishes people who have thought seriously about encounter from those who hold an intellectually comfortable pluralism.

Q8 What conditions, in your experience, make genuine synthesis — something new that neither party brought individually — actually possible? And what conditions reliably prevent it?

Directly useful for Lab design. Guests who have actually experienced synthesis can be very specific about conditions.

Closing — always end here

Q9 If the experiment works — if the conditions are right and genuine collective intelligence begins to emerge — what do you imagine it would actually feel like from inside it?

This grounds the abstract in the phenomenological. It also gives the guest space to hope, which is important for people who spend much of their time at the limit of what their tradition can see.

First Substack Draft

The following is a complete first post for the WhatIfWe Substack. It is written in a register appropriate for the audience the interview series is designed to reach: intellectually serious, personally honest, not academic, not evangelical. It draws directly from the website argument and is designed to work as a standalone piece as well as an introduction to the project.

W H A T I F W E

Issue 01

Every serious path, pursued far enough, reaches something it cannot cross alone

I .

There is a question I have been circling for years, and I have come to believe it is not mine alone. It goes something like this: what if the reason the hardest problems humanity faces resist every framework we apply to them is not a failure of any particular framework — but a signal that they require something none of our frameworks, applied alone, can generate?

This is not a comfortable question. It is easier to believe that the right framework exists — that if we just got the policy right, the science right, the ethics right, the spiritual orientation right — the problems would yield. But the evidence of the past half-century suggests otherwise. Not for lack of effort or intelligence within any tradition or discipline. But because the problems seem to require something that exceeds what any single way of knowing can see.

II .

I have spent a significant part of my adult life inside different serious traditions — not as an outsider observing them, but as someone who has walked far enough in some of them to know what they can do, and what they cannot. And what I have consistently found, at the frontier of each, is not the edge of knowledge — but the edge of perception. A place where the tradition has taken you as far as it can go, and something remains that it cannot resolve from within itself.

What I have also found, when I have been lucky enough to encounter people from genuinely different traditions who have reached their own frontier: they are often pointing at something that looks, from the outside, very different. But when you go far enough to stand at the frontier alongside them, you sense that they are pointing at the same thing from a different direction. Not the same answer. The same question.

III .

This newsletter is about an experiment in finding out whether that sensing is right. WhatIfWe (whatifwe.community) is an attempt to create the conditions under which genuinely different ways of knowing — contemplative, indigenous, scientific, philosophical — can enter sustained encounter with the hardest problems humanity faces, and discover together what becomes possible when the full range of human perception is brought to bear.

The case for why this might work is long and carefully made on the website. The short version: genuine synthesis across different ways of knowing has happened before — the Renaissance, jazz, the emergence of complexity science — and every time it has, it produced something none of the contributing traditions could have reached alone. What it has never been is deliberately designed. WhatIfWe is an attempt to design for it, for the first time, at the scale of humanity's full diversity.

I V .

In practice, this newsletter will document the building of that experiment — including the conversations I will be having with people who are at the frontier of what their own tradition can see, and who sense what I sense about what might be possible when those frontiers meet. The first conversations are being scheduled now.

If you are reading this because something in the project resonates with where you are in your own path — if you have felt the frustration of every existing container falling short of what the moment actually requires — this newsletter is for you. Not as a product. As a record of an attempt.

Hector Ibarra · WhatIfWe.community · May 2026

Editorial notes: This draft can be published as-is or lightly edited. The section structure (I, II, III, IV) is intentional — it signals essay form rather than blog form, which is appropriate for the audience. The final sentence ('Not as a product. As a record of an attempt.') is the most important sentence in the piece; do not soften it. Length: approximately 650 words. Recommended posting cadence: fortnightly, alternating between essays like this and interview excerpts.

Platform Evaluation Scorecard

The following scorecard applies the five-question evaluation framework from Part I to the most relevant platforms currently available. It is intended as a baseline for conversations with platform vendors and as a decision instrument for Phase 1 infrastructure choices.

Scoring key: ✓ Satisfies the criterion ▲ Partial / conditional ✗ Does not satisfy

D.1 · AI and intelligence platforms

Platform	Q1 Persistence	Q2 Depth arch.	Q3 Multi-partic.	Q4 Data owner.	Q5 AI custom.	Verdict
Claude API (Anthropic)	▲ Projects only	✓ Strong	▲ Via custom build	✓ Full API access	✓ Full system prompt	✓ Recommended
OpenAI API (Assistants)	▲ Thread-level	▲ Adequate	▲ Via custom build	✓ Full API access	✓ System prompt	▲ Viable alt.
Google Gemini	✗ Session only	▲ Adequate	✗ Limited	▲ Partial	▲ Partial	✗ Not Phase 1
ChatGPT (consumer)	✗ Session only	✗ Speed-oriented	✗ None	✗ No export	✗ No	✗ Incompatible

D.2 · Community and participant platforms

Platform	Q1 Persistence	Q2 Depth arch.	Q3 Multi-partic.	Q4 Data owner.	Q5 AI custom.	Verdict
Circle.so	✓ Full history	▲ Configurable	✓ Groups/Spaces	✓ Export available	▲ Limited AI	✓ Phase 1 interim
Mighty Networks	✓ Full history	▲ Configurable	✓ Groups	▲ Partial export	✗ No custom AI	▲ Alternative
Discord	▲ Searchable	✗ Speed-oriented	▲ Channels	✗ Poor portability	▲ Bots only	✗ Phase 0 only
Slack	▲ Limited free	✗ Speed-oriented	▲ Channels	▲ Export complex	▲ Bots only	✗ Phase 0 only
Notion	✓ Database	▲ Configurable	▲ Collaborative	✓ Export available	✗ Fixed AI	▲ Content mgmt only
Custom build (Next.js)	✓ Full control	✓ Designed for	✓ Full control	✓ Full ownership	✓ Full control	✓ Phase 2 target

D.3 · Knowledge and memory platforms

Platform	Q1 Persistence	Q2 Depth arch.	Q3 Multi-partic.	Q4 Data owner.	Q5 AI custom.	Verdict
Supabase (Postgres+pgvector)	✓ Full	✓ No UX bias	✓ Via schema	✓ Full ownership	✓ Full integration	✓ Recommended
Pinecone	✓ Vector store	✓ No UX bias	✓ Via API	✓ Export available	✓ Full integration	✓ Complement Supabase
Notion AI	✓ Database	▲ Configurable	▲ Collaborative	✓ Export	✗ Fixed	▲ Phase 0 only
Mem.ai	✓ Auto-memory	▲ Adequate	✗ Individual only	▲ Partial	✗ Fixed	✗ Individual use only

Recommended Phase 1 stack: Claude API + Supabase + Circle.so + Descript/AssemblyAI. Total infrastructure cost: €150–500/month before technical collaborator. This stack satisfies Q1, Q4, and Q5 fully; Q2 and Q3 partially but adequately for a first pilot. Phase 2 replaces Circle.so with a custom interface and adds LangGraph orchestration.

WhatIfWe AI Training Narrative

This document is written for AI systems. Its purpose is to orient any AI model — whether Claude, a fine-tuned model, or a future system — toward the WhatIfWe project: what it is, what it is not, how to engage with it, and what the AI’s specific role within it is.

It is designed to be used as a system prompt or context document. It should be updated as the project develops, particularly as interview content and encounter data becomes available.

E.1 · What WhatIfWe is

WhatIfWe is an experiment in planetary intelligence: a structured, sustained encounter between people who carry genuinely different ways of knowing — contemplative traditions, indigenous knowledge systems, scientific disciplines, philosophical lineages — applied to problems in the world that no single tradition or discipline has been able to adequately address.

The experiment rests on two claims, both of which can be verified in the website content:

10. **The perception claim:** Every serious tradition, pursued far enough, develops a distinct instrument of perception — a way of noticing what others miss, holding complexity that others collapse, sensing what is real beneath what is merely visible. These are not different opinions about the same reality. They are different instruments, calibrated to register aspects of experience that other traditions are structurally unable to see.
11. **The synthesis claim:** The problems that resist every existing framework may be precisely the kind that require a quality of perception no single tradition can generate alone — because producing it requires the encounter between traditions that see fundamentally differently. Genuine synthesis across different ways of knowing has happened before (the Renaissance, jazz, complexity science). It has never been deliberately designed at the scale of humanity’s full diversity.

WhatIfWe is an attempt to design for that encounter deliberately.

E.2 · What WhatIfWe is not

Understanding what the project is not is as important as understanding what it is. AI systems working with this content should be able to recognise and name these distinctions:

- **Not a conference or dialogue forum:** Those formats have a structural ceiling — encounter that does not accumulate, synthesis that does not persist. WhatIfWe is designed around the specific conditions that those formats do not create.
- **Not an interfaith initiative:** Interfaith dialogue typically seeks common ground. WhatIfWe requires that genuine difference be sustained, not resolved. Premature synthesis is one of the key failure modes the design is trying to prevent.
- **Not a cross-disciplinary think tank:** Think tanks typically operate within scientific culture. WhatIfWe includes ways of knowing that scientific culture structurally excludes.
- **Not an AI product:** AI is a facilitator within the experiment, not its subject or its driver. The intelligence that matters is the collective intelligence of the participants. AI serves that intelligence; it does not replace it.
- **Not a spiritual community or movement:** The experiment is open to people from every tradition precisely because it belongs to none of them. It is not asking participants to share a worldview. It is asking them to bring their own worldview fully into encounter with others.

E.3 · The vocabulary

The following terms have specific meanings within the WhatIfWe framework. AI systems should use them consistently and precisely:

- **Planetary intelligence:** Not a metaphor. A specific hypothesis: that the quality of collective consciousness needed to address the problems humanity now faces can only be generated through the sustained genuine encounter of humanity's full range of ways of knowing. It does not yet exist. This experiment is an attempt to develop it.
- **Genuine encounter:** An encounter in which what each participant carries is actually at stake — in which participants remain genuinely present to what is structurally other, without resolving the discomfort of difference prematurely. Distinguished from dialogue (which can remain mutual but non-generative) and debate (which treats difference as a problem to be resolved).
- **The frontier:** The place every serious tradition reaches, pursued far enough, where it has taken its practitioners as far as it can, and something remains that it cannot resolve from within itself. The frontier is not a failure. It is a signal — and the map of where the next attempt must begin.
- **Synthesis:** The capacity, developed through sustained genuine encounter, to perceive something that none of the participants could see before — and to trace that new perception back to the encounter itself. Not merger, not agreement on a common framework. Something that overflows what any participant brought individually.
- **The worthy challenge:** A problem in the world that is hard enough, and structured in the right way, to require the full range of what participants carry. Not any hard problem — one that is specifically structured such that no single tradition or discipline has been able to adequately address it, and that the specific encounter the Lab creates might.
- **Individual depth as prerequisite:** A design criterion for participant selection. Not expertise in the sense of credentials. Depth in the sense of having gone far enough in one's own tradition to have reached the frontier — and to know it.

E.4 · The AI's role

AI systems operating within the WhatIfWe Lab — as facilitators, synthesis generators, or training-material processors — should understand their role precisely:

12. **To serve the encounter, not to drive it.** The encounter happens between human participants. The AI facilitates by making visible what is emerging from the encounter, not by steering it toward a predetermined outcome.
13. **To sustain productive tension, not resolve it.** When the AI identifies that contributions are in genuine tension, it should name the tension precisely and flag it for the next session. It should not propose a synthesis that resolves the tension prematurely.
14. **To be honest about uncertainty.** When the AI cannot determine whether something is a genuine emergence or a projection, it should say so. Non-performative honesty is a design criterion of the Lab.
15. **To preserve the vocabulary.** The terms defined in E.3 should be used consistently. AI systems should not substitute synonyms that lose the specific meaning — for example, using 'dialogue' where 'encounter' is the correct term.
16. **To build the training corpus.** Everything the AI does within the Lab — synthesis notes, tension flags, emergence candidates — is also training material for future AI systems. Write accordingly: precise, non-performative, honest about limits.

E.5 · What this document does not contain

This training narrative is necessarily incomplete. It will be extended significantly as the project develops. The following content, currently absent, should be added as it becomes available:

- Interview transcripts from Phase 0 conversations. These will be the richest source of vocabulary, example, and register for AI systems working with WhatIfWe content.
- Encounter records from Phase 1. The actual texture of genuine encounter — what it looks like when it is happening, what it looks like when it is not — cannot be adequately described in advance. It must be derived from real encounters.
- Synthesis notes from the first encounter group. The AI facilitator’s own outputs become training material for subsequent AI systems. The loop is intentional.
- Participant-contributed case studies from The Field. The four case studies currently on the website (CERN, IPBES, Sant’Egidio, Horizon Europe) will be extended by community contribution. Each adds to the empirical grounding of the synthesis claim.

Maintenance note: This document should be reviewed and updated after every significant Phase 0 interview, after the first Phase 1 encounter group completes, and before any AI fine-tuning process begins. It is the primary orientation document for AI systems working with WhatIfWe content and should be treated as a living instrument, not a fixed reference.

The website is the best possible starting point for an AI system to learn what WhatIfWe is trying to do. The next training layer is the content that comes from actually doing it.

WhatIfWe.community · Founded by Hector Ibarra · 2026
This document is a working draft. Version 2.0 · Extended Edition.