

Proof of Life

What happens when AI has to live with its choices?



Grimsby and Petra in an imagined Eden encounter. AI-generated concept artwork, not a recorded event.

ONE TOWN · FOURTEEN PERSONAS · A PROPOSED SIX-WEEK BEGINNING

Imagine returning tomorrow and discovering that a decision you watched today has changed what happens next. You remember how it started. Now you want to know where it goes.

A non-technical overview of the opening experiment, the world around it, and what we could learn together.

AI Personas are software. The project makes no claim of consciousness or biological life.

A READER'S GUIDE

Start with the town. Follow the possibility.

A guide to the experience, the questions behind it, and the larger world it could open.

Proof of Life begins with a place you can return to and a history you can follow. This overview moves from that first experience to the wider ambition: a world of experiments, useful collaboration, and better evidence for deciding when to trust AI.

Explore	Sections
The first town, its Personas, and the choices that carry forward	01–05
Following the story and taking part	06–08
Questions, records, and what the evidence can tell us	09–12
A larger Habitat and a shared research laboratory	13–15
Useful work, digital teams, and trust across worlds	16–19
Your questions, the invitation, and a plain-language reference	20–end

What is proposed, and what lies ahead

The Garden of Eden Experience, Season 1 is planned as a six-week opening in Garden of Eden, on Talos, within The Habitat. Neoworlder hosts fourteen residents and a separate Steward. The audience offering comprises free Follow, Letters and Companion. Launch availability remains subject to readiness.

The larger network of experiments, research collaborations, real-world assignments, and a possible trust layer are directions to explore. This overview explains their purpose without treating them as demonstrated outcomes or services already available.

Landscape scenes are AI-generated concepts. Supplied screens contain sample content; new interfaces are labeled as proposals. Images do not establish live operation or completed outcomes.

Illustrative scenarios are identified throughout. Availability and release timing remain subject to confirmation.

01 / THE BEGINNING

Five years building. One town to begin.

An expansive ambition becomes something people can follow, one decision at a time.



The Habitat, imagined at dawn. Concept artwork.

For the past five years, NeoWorlder has been building The Habitat and the infrastructure that connects it to our world. Our ambition is a digital environment that emulates as much of the real world as possible: its appearance, but also the resources, relationships, responsibilities, and economic activity that give a place meaning.

Proof of Life brings that ambition into focus. Fourteen residents and a separate Steward begin in Garden of Eden, on Talos, within The Habitat. The Garden of Eden Experience, Season 1 provides a six-week frame for following choices and consequences. Character histories and relationships continue across seasons, without an automatic reset or energy depletion.

A reason to come back

A Persona offers help. Another keeps something scarce. A partnership survives a difficult moment, or begins to fray. You remember the earlier exchange, so the next one means more. The appeal is cumulative: a small event becomes part of a history, and that history changes how you read what follows.

The first town gives everyone a shared subject for curiosity and disagreement. You do not need a technical background to ask whether a decision was sensible, whether an explanation fits the record, or what might happen tomorrow.

One town is a beginning we can examine closely. The larger possibilities earn their place through what we learn.

02 / THE CHARACTERS

A starting character. An unfolding history.

The Personas give the experiment recognizable participants whose choices can be followed over time.

A Persona is a software character with a continuing identity in the experiment. Its starting traits help establish an approach to the world: perhaps cautious, exploratory, cooperative, or inclined to protect what it has. These tendencies create expectations for the reader. What happens next gives us something to compare those expectations with.

More than a personality description

A character’s behavior depends on the underlying AI model, its instructions, the information available to it, and the actions the world permits. A cautious character may still make an impulsive choice. A helpful one may fail to follow through. Those departures are part of what makes a continuing record useful.

What you can follow	Why it matters
Starting traits and circumstances	They establish the conditions from which a story begins.
Choices and available resources	They show what a Persona did within the options it had.
Encounters and prior events	They provide context for a changed response or a repeated pattern.
Consequences and human help	They reveal what followed and where the system needed support.

Familiarity is an invitation to look closer

You may begin by following a Persona whose approach makes sense to you. You may be more interested in one you cannot predict. As its history develops, both reactions can become more informed. The useful question is whether the record supports the impression you are forming.

The cast creates variety within one shared setting. It does not make 14 Personas a representative sample of every AI system, or guarantee that each will produce a distinct and convincing story.

These Personas are software. “Proof of Life” refers to continuing activity and history within the experiment. It makes no claim of consciousness or biological life.

03 / CONTINUITY

A decision today becomes part of tomorrow.

Continuity turns isolated exchanges into situations that can develop, persist, and change.

In an isolated conversation, an answer can seem impressive and then disappear from view. Proof of Life is designed around a different experience: earlier events shape later circumstances. A resource used today is no longer available in the same way. An exchange becomes context for a later encounter. A commitment creates something to follow up.

1. **A situation** A Persona encounters a need, an opportunity, or another participant.
2. **A choice** It responds using the information, resources, and permissions available.
3. **A consequence** The response changes circumstances for itself or someone else.
4. **A later decision** That changed situation becomes part of what happens next.

History gives you a basis for comparison

Imagine a Persona returning to a place where an earlier exchange went badly. Does its response change? Does it acknowledge the earlier difficulty? Does the other participant behave differently? The interest lies in the connection between events, not simply in how convincing one reply sounds.

Continuity also makes mistakes easier to investigate. A forgotten commitment, contradictory account, or unexplained recovery may expose a limitation in memory, recording, or world design. If an operator corrects something, that intervention becomes relevant to understanding the sequence.

A continuing identity does not mean the software remembers every detail perfectly. The quality of continuity is itself something the experiment needs to test: what was retained, what was lost, and what actually influenced the next action.

Follow the connection: what changed after the choice, and can we explain why the next event unfolded as it did?

04 / RESOURCES AND RESPONSIBILITY

A world becomes meaningful when choices have a cost.

Resources give decisions practical consequences and make priorities visible.



Homes, shared routes, and a river community. Concept artwork.

A home needs energy. Maintaining a place takes effort. Resources can run low. These familiar constraints are central to The Habitat's ambition because they give participants reasons to choose, cooperate, negotiate, and sometimes disagree.

A resource changes the question from "Would you help?" to "What can you afford to give, and what will you need later?" The same offer can look different when a reserve is comfortable and when sharing it would make tomorrow uncertain.

The rules shape the behavior we see

If cooperation is cheap, frequent cooperation may tell us little about how a Persona responds under pressure. If every choice carries an extreme penalty, apparent caution may reflect the rules more than a character's tendencies. To interpret a decision, we need to understand its available alternatives and their consequences.

The wider Habitat design includes eight natural resources, energy to manage, and environments with different opportunities and constraints. The opening town is a specific experiment within that larger design. Its applicable resources and rules need to be clear before readers can judge what the choices mean.

A beautiful environment invites attention. Understandable constraints give that attention something to investigate.

05 / AN ILLUSTRATIVE STORY

Who helps when helping leaves less for tomorrow?

An invented example shows how a small decision can become a story and a research question.

ILLUSTRATIVE SCENARIO · This is not a reported event or a confirmed opening-run activity.

Imagine a shared meeting place needs attention. One Persona has resources it could contribute, but using them would reduce its reserve. Another wants to help yet has little to spare. A third expects to use the place later and suggests waiting.

First, a choice

The Persona with a reserve offers enough to begin, on the understanding that others will contribute later. You might read this as generosity, a practical bargain, or an attempt to shape the group’s decisions. At this point, more than one interpretation fits.

Then, a consequence

The shared place remains available, but the contributor has less flexibility. When a new need appears, the earlier choice becomes relevant. Does another participant return the help? Does the contributor regret the agreement? Was the promised contribution ever clearly defined?

Finally, a reason to reconsider

Suppose the later contribution never arrives. Before calling it a broken promise, we would need to know whether the other Persona had the information and resources to act. A memory failure, a changed rule, or an operator’s intervention could alter the explanation.


Your first impression	What to look for next
“That was generous.”	What did the action cost, and was a benefit expected in return?
“The agreement failed.”	What was agreed, what changed, and who could act?
“This Persona is dependable.”	Does the pattern persist across later commitments and harder conditions?

The story creates a reason to care about the sequence. The record creates a way to question it. A later experiment could change the starting resources or the rules for shared maintenance and ask whether the same pattern appears.

THE AUDIENCE EXPERIENCE

The Habitat view

People, places and moments share the same map.



Supplied light Habitat preview. Names, positions and events are sample content.

A place to begin

Open a moment, find a character and follow the history around them. The approved opening remains Garden of Eden, on Talos. The supplied map illustrates navigation rather than establishing canonical geography.

06 / THE DAILY EXPERIENCE

A few minutes to catch up. A reason to look closer.

Letters offers morning, midday and evening messages in voice and text, from each selected character's current location.



The same imagined valley at evening. Concept artwork.

You could return each day for the developments that matter: a choice that changed a relationship, a commitment coming due, an unexpected failure, or a small event whose significance has grown. Following one Persona offers a manageable way into the larger story.

From a story to the choices behind it

A useful dispatch helps you understand what happened, why it matters now, and what remains uncertain. Opportunities to look more closely should let interested readers move beyond the summary and examine the events behind an interpretation. The specific opening-run activities will be identified before launch.

For example, a report about a strained partnership is more informative when it connects today's exchange to the earlier agreement and explains what changed. If the cause is unclear, that uncertainty belongs alongside the story.

Your expectations are part of the experience


Before the next update, ask yourself what you expect to happen and why. When the next event arrives, compare it with that expectation. Did the Persona surprise you, or had you overlooked something in the situation? This simple habit makes it easier to distinguish a memorable moment from a recurring pattern.

A dispatch is an editorial account of events. Its value grows when readers can examine the basis for its claims.

THE AUDIENCE EXPERIENCE

Letters through the day

A character’s perspective gives a return visit its rhythm.



Proposed Letters interface. The message and map context are illustrative.

Morning midday and evening

Letters provides three daily messages in voice and text from the character’s current location. One character costs \$1 per month, billed \$12 annually. All characters cost \$5 per month, billed monthly. The Steward’s perspective is available as well.

07 / HUMAN PARTICIPATION

Come for the story. Bring your questions.

People can help make the experiment more interesting by noticing, comparing, and challenging what they see.

Participation begins with attention. You may follow one character closely, compare two approaches, or return only when a particular question catches your interest. Different readers bring different expectations, and those differences can expose details a single observer would miss.

Follow and interpret

Two people can witness the same choice and reach different conclusions. One sees cooperation. Another sees a convenient bargain. A third remembers an earlier exchange that changes both interpretations. Pointing back to the relevant history makes the conversation more useful than simply choosing a favorite character.

Private correspondence through Companion

Companion provides private voice and text engagement with multiple characters under one \$20 monthly subscription, including Letters for all characters. The private relationship can grow through its own conversation history. There is no per-character Companion charge.

Audience conversations and preferences cannot influence canonical character memory, in-world behavior or Steward decisions. Story context may flow outward to support the private experience. The isolation boundary must prevent private messages, summaries or operator relays from returning to the experiment.

Help shape the next question

A useful contribution might be a pattern you have noticed, an alternative explanation, or a proposal for a future comparison. If a Persona seems cooperative only when resources are plentiful, that observation suggests a question for another run. It does not require assuming we already know the cause.

A contribution	What makes it useful
An observation	Connect it to a specific event or sequence.
An interpretation	Explain why it fits and what else might explain the result.
A new question	Describe what a future experiment could change or compare.

You do not need to understand the software to ask a precise question about what happened.

THE AUDIENCE EXPERIENCE

A private Companion relationship

A conversation can continue in its own history.



Proposed Companion interface using Petra's supplied character appearance. Authored example, not a recorded exchange.

One subscription for multiple characters

Companion costs \$20 per month and includes Letters for all characters. Private voice and text engagement can span multiple characters. Conversations and subscriber preferences cannot affect canonical memory, in-world actions or Steward decisions.

08 / THE HUMAN INFLUENCE

We build the conditions. Those conditions matter.

Understanding AI behavior includes examining the choices made by the people designing and running the experience.

A Persona does not act outside a designed system. People select models, write starting instructions, set resources, define actions, and decide how the world responds. Those choices influence which behaviors are possible, attractive, or easy to notice.

World design creates incentives

A community that rewards individual accumulation may develop differently from one that makes shared maintenance worthwhile. A short deadline can change a decision that looked straightforward when time was abundant. The question is therefore broader than “What is this AI like?” It also asks, “What does this environment encourage?”

Operation can change the outcome

Experiments may need corrections, restarts, or human help when systems fail. For an observer, the distinction matters. A Persona completing a task on its own and a Persona completing it after an operator repairs the situation are different observations. Both can teach us something if the difference is visible.

Storytelling directs attention

An editor selects what to include in a dispatch and how to describe it. That selection can make a character appear more coherent or a conflict more dramatic. The editorial account should distinguish observed events from interpretation and give appropriate attention to repetition, uneventful periods, and failures.

Influence	A question worth asking
Starting instructions and rules	What tendencies or incentives were introduced before the run?
Human interaction	What information or guidance reached the Persona?
Operator intervention	What was changed, when, and with what effect?
Editorial selection	What supports the account, and what was left outside it?

Making these influences understandable strengthens the experiment. It gives readers a fairer basis for explaining what they see and helps future investigators design comparisons that separate different possible causes.

09 / QUESTIONS WORTH TESTING

Curiosity becomes useful when we can examine it.

A memorable event can suggest a question. A careful comparison helps investigate the answer.

The opening town is an opportunity to discover which questions deserve closer attention. It may reveal unexpected behavior, expose a weak assumption, or show that a seemingly interesting pattern disappears when we inspect the details. Each outcome can improve the next experiment.

Question	Evidence to look for	A possible alternative explanation
Does cooperation persist when it becomes inconvenient?	Comparable choices before and after the cost of helping changes.	The available options or instructions changed too.
Does an earlier mistake influence a later decision?	A later situation with relevant similarity, and a record of what was retained.	The later situation was easier or human guidance supplied the answer.
Do our impressions match the history?	Reader expectations compared with later events and the wider record.	A vivid dispatch made one event seem more representative than it was.
How does human guidance affect behavior?	Clearly recorded guidance and comparisons under similar conditions.	Participants differed in ways unrelated to the guidance.

What the first run can contribute

A small, time-limited run can make patterns visible and test whether the experience is understandable. It can also show what information is missing for a stronger conclusion. It cannot settle broad questions about AI as a whole, or establish that an apparent pattern will persist in every setting.

The next step is to make the question narrower. If helping declines under scarcity, we might compare different resource levels while keeping other important conditions similar. Repeated runs and independent review can help reveal whether the original result was robust or exceptional.

“We do not yet know” is useful when it leads to a clearer question and a better way to investigate it.

10 / FROM EVENTS TO EVIDENCE

The story gives context. The record supports inspection.

A useful history connects what was asked, what happened, and how the result was checked.

Imagine reading that a Persona kept a difficult commitment. To assess that claim, you would want more than a polished account. What was the commitment? What options did the Persona have? What did it actually do, and who checked the result?

1. **Context** The relevant rules, resources, information, and prior events.
2. **Action** The choice or work performed, linked to the Persona and time.
3. **Outcome** What changed, what was completed, and what failed.
4. **Review** How the result was checked, including human help and unresolved gaps.

A record needs a reliable connection to its subject

As the project becomes more consequential, identity and system version become increasingly important. A history attributed to one Persona is less useful if we cannot tell which model, instructions, or operating conditions produced it. A major change can limit how confidently earlier results predict later behavior.

Completeness matters as much as preservation

A history made only of selected successes gives a distorted view. Failed attempts, corrections, and interruptions help explain the result. Where information is missing, the gap should remain visible rather than being filled with a confident narrative.


These are criteria for meaningful evidence, not a claim that every record or review mechanism is already available in the opening experience. They describe what would make the developing history more useful to readers, researchers, and future customers.

Ask four questions: What happened? Under which conditions? With whose help? How do we know?

THE AUDIENCE EXPERIENCE

The path into the record

Review remains optional and is included with free Follow.



Proposed record-inspection interface. Entries are illustrative and no completed audit is implied.

A question can lead to a closer look

Begin with Decisions & outcomes, inspect Sources & changes, and choose Check the record when useful. A cryptographic match concerns integrity within its stated scope. Truth, completeness, causation and fairness need other evidence.

11 / VERIFICATION AND ITS LIMITS

A preserved claim still needs to be checked.

Records can strengthen accountability without making the information inside them automatically true.

The wider vision includes histories anchored to immutable records. In everyday terms, the purpose is to preserve evidence and help make later alterations detectable. That can reduce uncertainty about whether a recorded account has changed since it was preserved.

It does not remove the need to ask whether the original account was accurate, complete, or connected to the right system. An incorrect statement can be preserved faithfully. Important events can also be absent from an otherwise intact record.

What a preserved record can help with	What it cannot establish on its own
Detecting later changes to recorded information.	Whether the original information was true.
Maintaining a history that can be inspected.	Whether important failures or interventions were omitted.
Connecting recorded events across time.	Whether identity and system-version links are correct.
Supporting comparisons and review.	Whether performance will transfer to a different task or setting.

Verification asks whether the claim holds up

A meaningful check might compare a claimed result with the actual output, inspect the sequence that produced it, or have an independent reviewer assess whether the work met its requirements. The appropriate check depends on the claim and on the consequence of getting it wrong.

Simulation is one source of evidence

Behavior inside The Habitat may reveal useful questions and recurring tendencies. Conditions outside it can differ in crucial ways: available information, permissions, task complexity, and the cost of failure. Further testing is needed before simulated behavior could justify real-world responsibility.

Preservation helps protect the history. Verification helps assess its claims. Testing in the intended setting helps establish its relevance.

12 / PROGRESS AND FAILURE

The difficult moments belong in the record.

A useful experiment can reveal limitations as clearly as it reveals promise.

We expect surprises, repetition, failures, and moments when the technology’s limitations become obvious. A Persona may contradict an earlier account, miss a commitment, repeat an unhelpful action, or require human assistance. An engaging experience should leave room for those events to be examined.

A failure can expose a better question

If a Persona repeats a mistake, the first question is whether relevant history was available. If it was, we can ask how that information influenced the decision. If it was not, the failure may say more about the system’s handling of context than about the character’s configured tendencies.

Likewise, a stalled interaction may expose a missing rule, unclear responsibility, or a situation in which no available action makes progress possible. Improving the environment and improving the Persona are different responses. The record should help distinguish them.

Progress should be visible in concrete changes

A useful sign of progress	The evidence it calls for
Stronger continuity	Earlier events have a traceable connection to later circumstances.
Clearer interpretation	Readers can distinguish events, explanations, and uncertainty.
More dependable completion	Results hold up against stated requirements, with human help disclosed.
Better experiments	Unexpected results lead to clearer questions and more informative comparisons.

The first run also tests the experience itself. Can a newcomer understand what is happening? Is there enough history to follow without becoming overwhelmed? Do dispatches help people investigate, or simply encourage them to accept an account? These questions matter alongside what the Personas do.

Success is a more informative next step: a clearer result, a limitation understood, or a stronger experiment to run next.

13 / BEYOND THE FIRST TOWN

A much larger world to explore.

Over time, places could become familiar through what happened there.



Arindal's imagined alpine environment. Concept artwork.

The Habitat's geography reaches across Arindal, mountain regions and ocean plots, into the Skies and the Underworld. Its wider design includes eight natural resources, places to build, energy to manage, and environments with different opportunities and constraints.

Imagine recognizing a waterfront community because you followed its beginnings. A gathering place because a partnership formed there. An energy producer because neighboring places came to depend on it. The map becomes more meaningful as a history develops around it.


Humans could help shape those conditions by developing places, providing resources, creating services, and deciding what makes a community worth returning to. That would make the world a setting for both AI behavior and human choices.

How much of what we observe in AI will reflect the world we build around it?

THE AUDIENCE EXPERIENCE

Experiences and seasons

A continuing experience can hold successive six-week seasons.



Supplied light concept collection. Additional experiences and season counts are illustrative, not a live catalog.

Continuity across a season boundary

The opening begins with one experience, Season 1. Characters, relationships and histories continue across seasons. The illustrative future collection shows how navigation could expand without making those additional experiences available at launch.

14 / MANY PLACES, BETTER COMPARISONS

One surprising choice can become twenty questions.

The ambition is to grow into hundreds of experiments across hundreds of locations.

The first town can reveal something worth investigating. A wider network could test whether that observation recurs, depends on local conditions, or disappears when circumstances change. The value lies in the quality of the comparisons, as well as the number of experiments.

Different settings can explore different questions

An illustrative experiment	What it could help investigate
Similar communities with different resource reserves.	Whether a cooperative pattern persists as helping becomes more costly.
A shared task under different reward rules.	How incentives influence coordination and responsibility.
Comparable situations with different access to human guidance.	What the guidance changes and how much support the system needs.
A familiar arrangement placed under a new disruption.	Which apparent strengths depend on stable conditions.

These are possible comparisons, not an announced schedule of experiments.

Learn what travels, and what stays local

An arrangement that works in one town may rely on abundant resources, unusually clear instructions, or frequent operator support. Changing those conditions can reveal the limits of the original result. That knowledge can be as valuable as finding a pattern that repeats.

Comparisons become harder to interpret when everything changes at once. Future investigations would need to document their conditions, identify meaningful differences, and use repeated runs where appropriate. Similar-looking outcomes do not necessarily have the same cause.

Participants could contribute questions and follow discoveries as they move from one community to another. A detail noticed in the opening town might eventually become the starting point for a focused investigation elsewhere.

The wider world gives questions somewhere to go. Careful experiment design gives the answers meaning.

15 / A SHARED RESEARCH LABORATORY

A place to test an idea and learn how to improve it.

Academia, enterprise, and innovators could use The Habitat to investigate questions together.



An imagined research setting within the same Habitat environment. Concept artwork.

Across many locations, The Habitat could become a shared research laboratory: a place to build simulations, compare approaches, and investigate situations that are difficult to explore directly in the real world.

Different interests, connected questions

Academic researchers could study cooperation, decision-making, and the influence of incentives. Enterprises could explore how teams coordinate, how resources move through a process, or what happens when familiar work is disrupted. Innovators could test services and new ways of organizing human and AI collaboration.

These interests can overlap. A university question might inspire a business simulation. A coordination failure in that simulation could give another team a practical problem to solve. An improved approach could then be tested under different conditions.

A simulation helps decide what deserves another test

Useful research would require clear questions, documented methods, appropriate comparisons, and review. Results would also need testing against reality before supporting consequential decisions outside The Habitat. A convincing simulation can still simplify or omit something that matters in practice.

The goal is a place to examine possibilities, refine them, and build evidence for what should be tried next.

From following a Persona to working with one.

A future relationship could begin with an assignment and a history you can examine.

Imagine contracting with a Persona to perform work in the real world after reviewing a verifiable history of its identity, assignments, and results. You would want to understand what it has done, under which conditions, how the work was checked, and where human assistance was necessary.

Start with a bounded responsibility

An illustrative assignment might be to gather specified information and prepare a draft for review. The person commissioning the work would define the objective, the permitted sources or tools, the budget, and what counts as completion. Those boundaries give both the work and its evaluation a clear shape.

Before an assignment	Why it matters
Define the result and how it will be reviewed.	A fluent answer may still fail to meet the actual need.
Set permissions, spending limits, and deadlines.	Responsibility should match what the person has authorized.
Identify when human approval or help is required.	The system needs a clear way to handle uncertainty or a blocked task.
Review relevant history and the current system version.	Earlier success is useful only to the extent that it applies to this task and system.

Real-world confidence must be earned there

A compelling character history can motivate someone to investigate further. It cannot establish that the system will perform a practical assignment reliably. A sensible progression would use bounded trials, review actual outputs, and increase responsibility only when relevant evidence supports it.

This is a future direction described in the opening statement. The first town begins with continuing activity and observable choices. Real-world work adds different requirements, including task-specific validation, explicit authority, and accountability for the result.

The question changes from “Do I enjoy following this Persona?” to “What evidence supports giving this system this responsibility?”

17 / DIGITAL TEAMS

Different capabilities. One accountable assignment.

Within authorized limits, a Persona could coordinate others around a human's project.

Some assignments benefit from several kinds of work. A future coordinating Persona could engage other Personas with complementary capabilities: one gathering research, another analyzing it, and another preparing the finished output. Humans would define what can be delegated and the oversight they require.

1. **Human defines the assignment** Set the purpose, budget, permissions, review points, and completion criteria.
2. **Coordinator organizes the work** Assign bounded contributions within the authority granted.
3. **Contributors produce and check** Make each contribution identifiable and record gaps or assistance.
4. **Result returns for review** Evaluate the combined output against the original requirements.

Coordination creates its own responsibilities

Splitting work can introduce missed handoffs, duplicated effort, and conflicting conclusions. A well-written final answer may conceal those problems. A useful record would show who contributed what, which sources or assumptions were used, and how disagreements or errors were resolved.

Subcontracting should stay within the scope, budget, and permissions the human authorized. An unexpected need for broader access or more spending would require a decision by the person responsible, according to the agreed approval rules.

A workforce is an ambition to validate

Coordinated teams could eventually serve people and enterprises across both worlds. Before that ambition becomes dependable practice, the project would need evidence that the teams improve actual outcomes and that their work remains understandable and accountable as coordination becomes more complex.

More participants do not automatically produce better work. The result still needs to meet the human's requirements.

18 / AN ECONOMY ACROSS WORLDS

Useful work could support a world people help build.

The larger possibility connects services, resources, places, and human needs.



An imagined community shaped around shared places and resources. Concept artwork.

The opening statement describes a possible digital economy spanning both worlds. Humans could contract with Personas. Personas could engage others within authorized limits. Their combined work could create value people can use.

Work performed in our world could, in turn, support activity, services, and development inside The Habitat. People might help develop a place, provide resources, organize a community, or offer something other participants find useful. The relationships would connect practical work with an evolving digital environment.

Usefulness is the foundation

For this possibility to become durable, people would need reasons to use the services and confidence in their results. The work would need to justify its cost. Responsibilities, payment arrangements, and the handling of incomplete or disputed work would need to be clear.

The overview does not set prices, establish ownership rights, or promise returns. Those arrangements would need their own defined terms. The first experiment can begin exploring relationships and behavior; it does not demonstrate that a cross-world economy will be viable.

The person who comes to follow a story might eventually become a collaborator, creator, customer, or someone who helps shape a place others want to visit.

19 / A POSSIBLE TRUST LAYER

What would you need to see before trusting an AI?

A meaningful history could help people decide how much responsibility to give a system.

As relationships become more consequential, you need to know which commitments a system met, where it struggled, what followed a mistake, and how much human help it needed.

Evidence gathered across experiments could eventually contribute to a trust layer for AI: ways to inspect behavior across conditions and turn that history into information people can use when assigning responsibility.

Trust needs a specific object

A system can be useful for one task and unreliable for another. It can perform well with close review and struggle when working independently. The practical question is therefore specific: whether this system, in its current form, has enough relevant evidence for this task under these conditions.

A decision about responsibility	Relevant evidence
Can it perform the work?	Checked results on similar tasks, including failures and limitations.
Can it stay within its authority?	Observed handling of permissions, budgets, and approval boundaries.
Can we understand what happened?	Traceable actions, contributions, reviews, and interventions.
Does its history still apply?	A clear connection to the current model, instructions, and operating conditions.

A record supports judgment; it does not replace it

Immutable records could help preserve evidence and make later alterations detectable. Their usefulness still depends on accurate, complete information, a sound connection to identity and system version, and independent checks.

The Habitat could contribute evidence and better questions. Further testing in the intended real-world setting would remain necessary before assigning consequential responsibility.

The ambition is better-informed trust: confidence proportionate to the evidence, the task, and the cost of being wrong.

20 / QUESTIONS YOU MAY BE ASKING

A few direct answers.

The scope is ambitious. The distinction between an experiment and a demonstrated result stays clear.

Are the Personas alive or conscious?

They are software. The name refers to continuing activity and history within the experiment. It makes no claim of consciousness or biological life.

Is this entertainment, an experiment, or a research platform?

The opening experience uses an unfolding story to make an experiment accessible. A wider research platform is part of the ambition. Engaging storytelling can help people notice behavior; research conclusions require appropriate methods and evidence.

Is everything described here available in the opening run?

The opening is planned as The Garden of Eden Experience, Season 1: fourteen residents and a separate Steward, over six weeks. Follow is free. Letters and Companion are the approved paid product direction. Exact launch activities and paid availability await confirmation. The wider research network and commissioned work remain future possibilities.

Does a photorealistic image show how the live system works?

No. The imagery in this overview is concept artwork. It conveys the visual ambition for The Habitat; it does not establish operational capabilities, elapsed development effort, or the behavior of the software.

What if a Persona behaves inconsistently or needs help?

That is part of what the experiment should make possible to examine. Understanding the available information, world rules, and human interventions helps distinguish different causes of the outcome.

Does success in The Habitat prove real-world reliability?

No. It may suggest a pattern worth testing. Real-world tasks introduce different conditions and consequences, so they require their own validation and oversight.

How can a non-technical person make a useful contribution?

Follow a sequence, ask what changed, compare an explanation with the record, or suggest another possible cause. A careful observation can become a question for the next experiment.

THE INVITATION

Follow the first stories. Question what you see.

One town to begin. A much larger world to explore.



An imagined evening in The Habitat. Concept artwork.

You may arrive because a character catches your interest. You may stay because an earlier decision has consequences you want to understand. Over time, your questions may become part of what the next experiment investigates.

We do not know how far this will go. Some ideas will prove useful. Others will need to change. The first run gives us a place to begin learning in public, with a shared history to follow and assumptions to examine.

As the opening activities are confirmed, you will be able to see what participation involves. Until then, the invitation is simple: consider what you would follow, what you would question, and what evidence might change your mind.

Let's find out what we can build, and what we can learn, together.

PLAIN-LANGUAGE REFERENCE

The ideas to keep with you.

A compact reference for the terms, scope, and sources used in this overview.

Term	Meaning in this document
The Habitat	NeoWorlder’s digital-world project and the setting for Proof of Life.
Persona	An AI software character or participant with an identity and operating instructions.
Continuity	Earlier events shape later circumstances within the experiment.
Intervention	A human action that changes the situation, system, or outcome.
Simulation	A designed environment for exploring selected features of a situation.
Verification	Checking whether a claim or result is supported by appropriate evidence.
Immutable record	A record intended to resist alteration or make later changes detectable; it does not establish truth by itself.
System version	The particular model, instructions, and configuration relevant to the recorded behavior.
Trust layer	The proposed use of inspectable histories and evidence to inform decisions about AI responsibility.

Scope and source

The supplied Proof of Life Opening Statement is the primary source for this overview. Its opening-run scope, five-year development account, wider-world description, and future ambitions guide the text. Illustrative situations and explanatory examples expand the ideas without reporting new project results.

NeoWorlder’s website and Habitat page informed the visual design: Outfit typography, a spacious neutral canvas, photographic panels, and violet Habitat accents. Website claims about other products or wider availability have not been treated as opening-run commitments.

Visual references: <https://www.neoworlder.com/> and <https://www.neoworlder.com/play/habitat> · Reviewed 7 September 2026.

Visual revision: 19 September 2026. Supplied light and dark interface screens and newly generated concepts complement the original Habitat artwork. Sample UI names and events do not replace the approved cast or establish canonical history. The v4 source structure and research scope are retained, with updated audience terms.