UNDERGROVE

Intelligence runs deeper.

A private AI workspace in development

Product brief / October 2026



The first audience

Independent builders and researchers

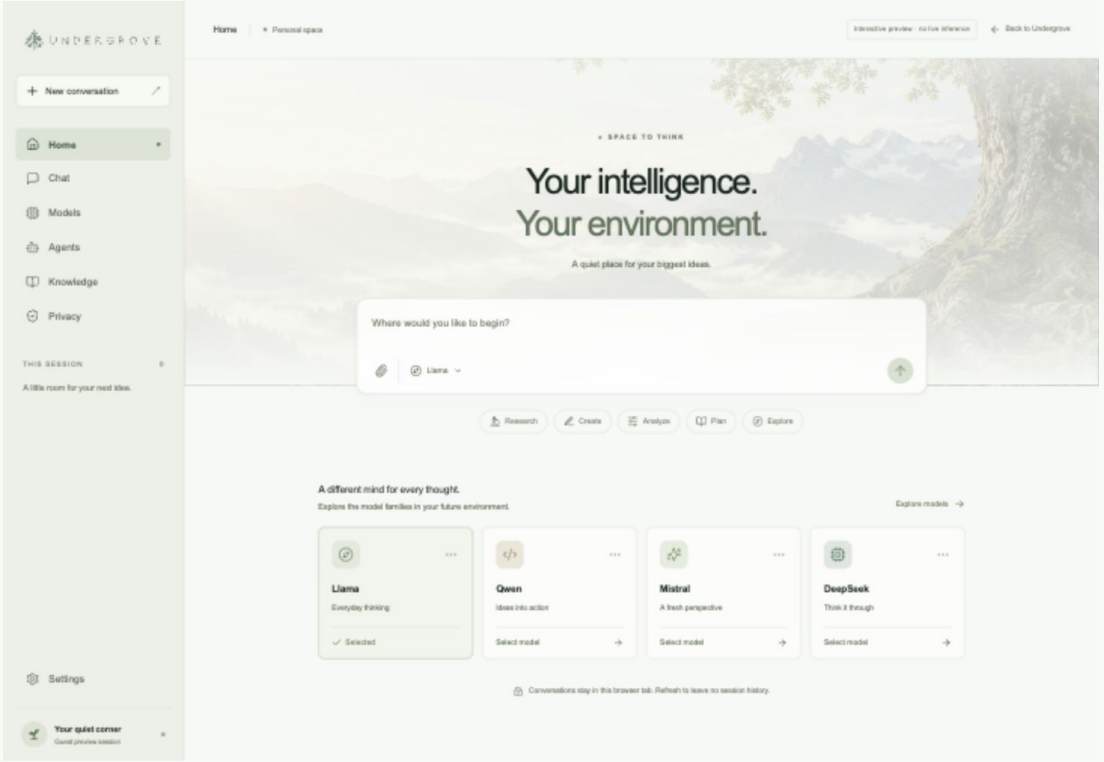
People who use AI for ongoing work and want to control the context they share.

STARTING HYPOTHESIS

A useful workspace should help them finish the task and understand where their information goes.



A workspace for sustained work



Current interface preview. Responses are scripted.

Conversation

Prompts, model selection and session history in one considered interface.

Personal knowledge

Planned retrieval and optional memory, with export and deletion controls.

What exists today

Built

Brand system and cinematic website

Responsive workspace

Scripted conversations in tab memory

File names, sizes and types only

Next to prove

Live inference and authentication

Storage, retrieval and deletion

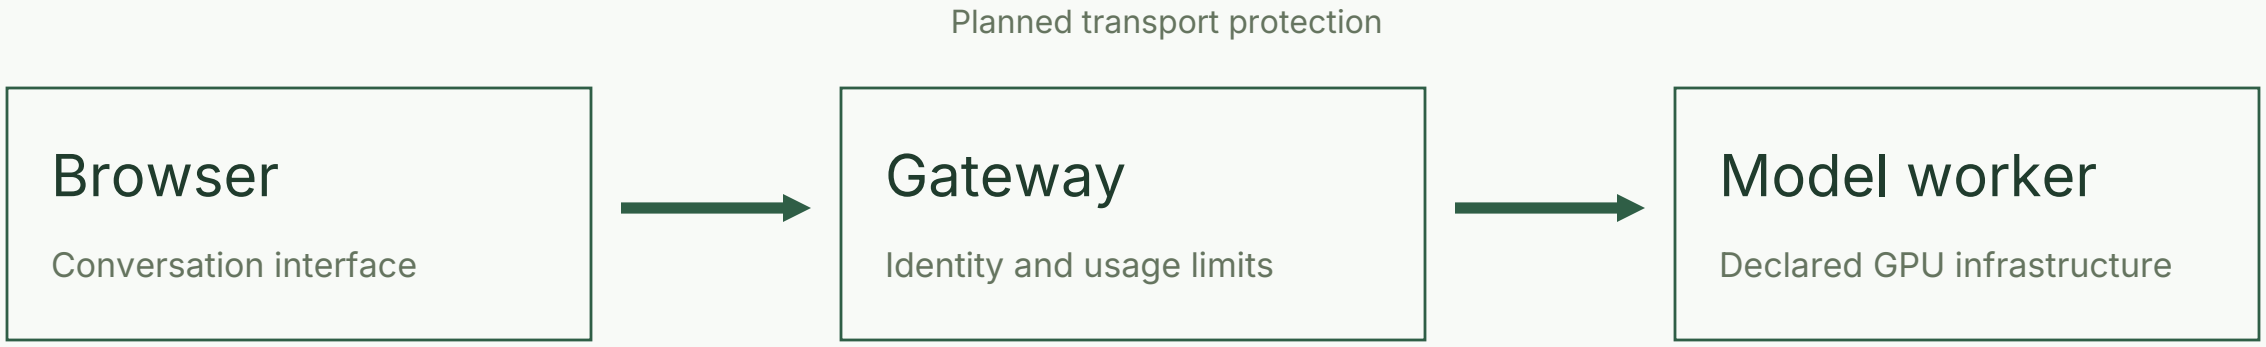
Retention and access controls

Costs, capacity and support

The preview does not run a model or store user accounts.

Proposed alpha data flow

One licensed open-weight model on a declared hosting route.



Plaintext boundary

The gateway and worker can process content in this first design.

Evidence before launch

A data inventory, retention tests and clear disclosures for each model route.

Our proposed position

Venice

PUBLISHED OFFERING

Model-dependent privacy modes

Anonymous
Private
TEE
E2EE

Undergrove

PRODUCT DIRECTION

A focused environment for sustained work

User-controlled knowledge
Clear data choices for each model
Portable conversations

Venice's published documentation, checked 19 October 2026. Provider claims are not an independent audit.

Release gates

Proposed sequence for the first live service.

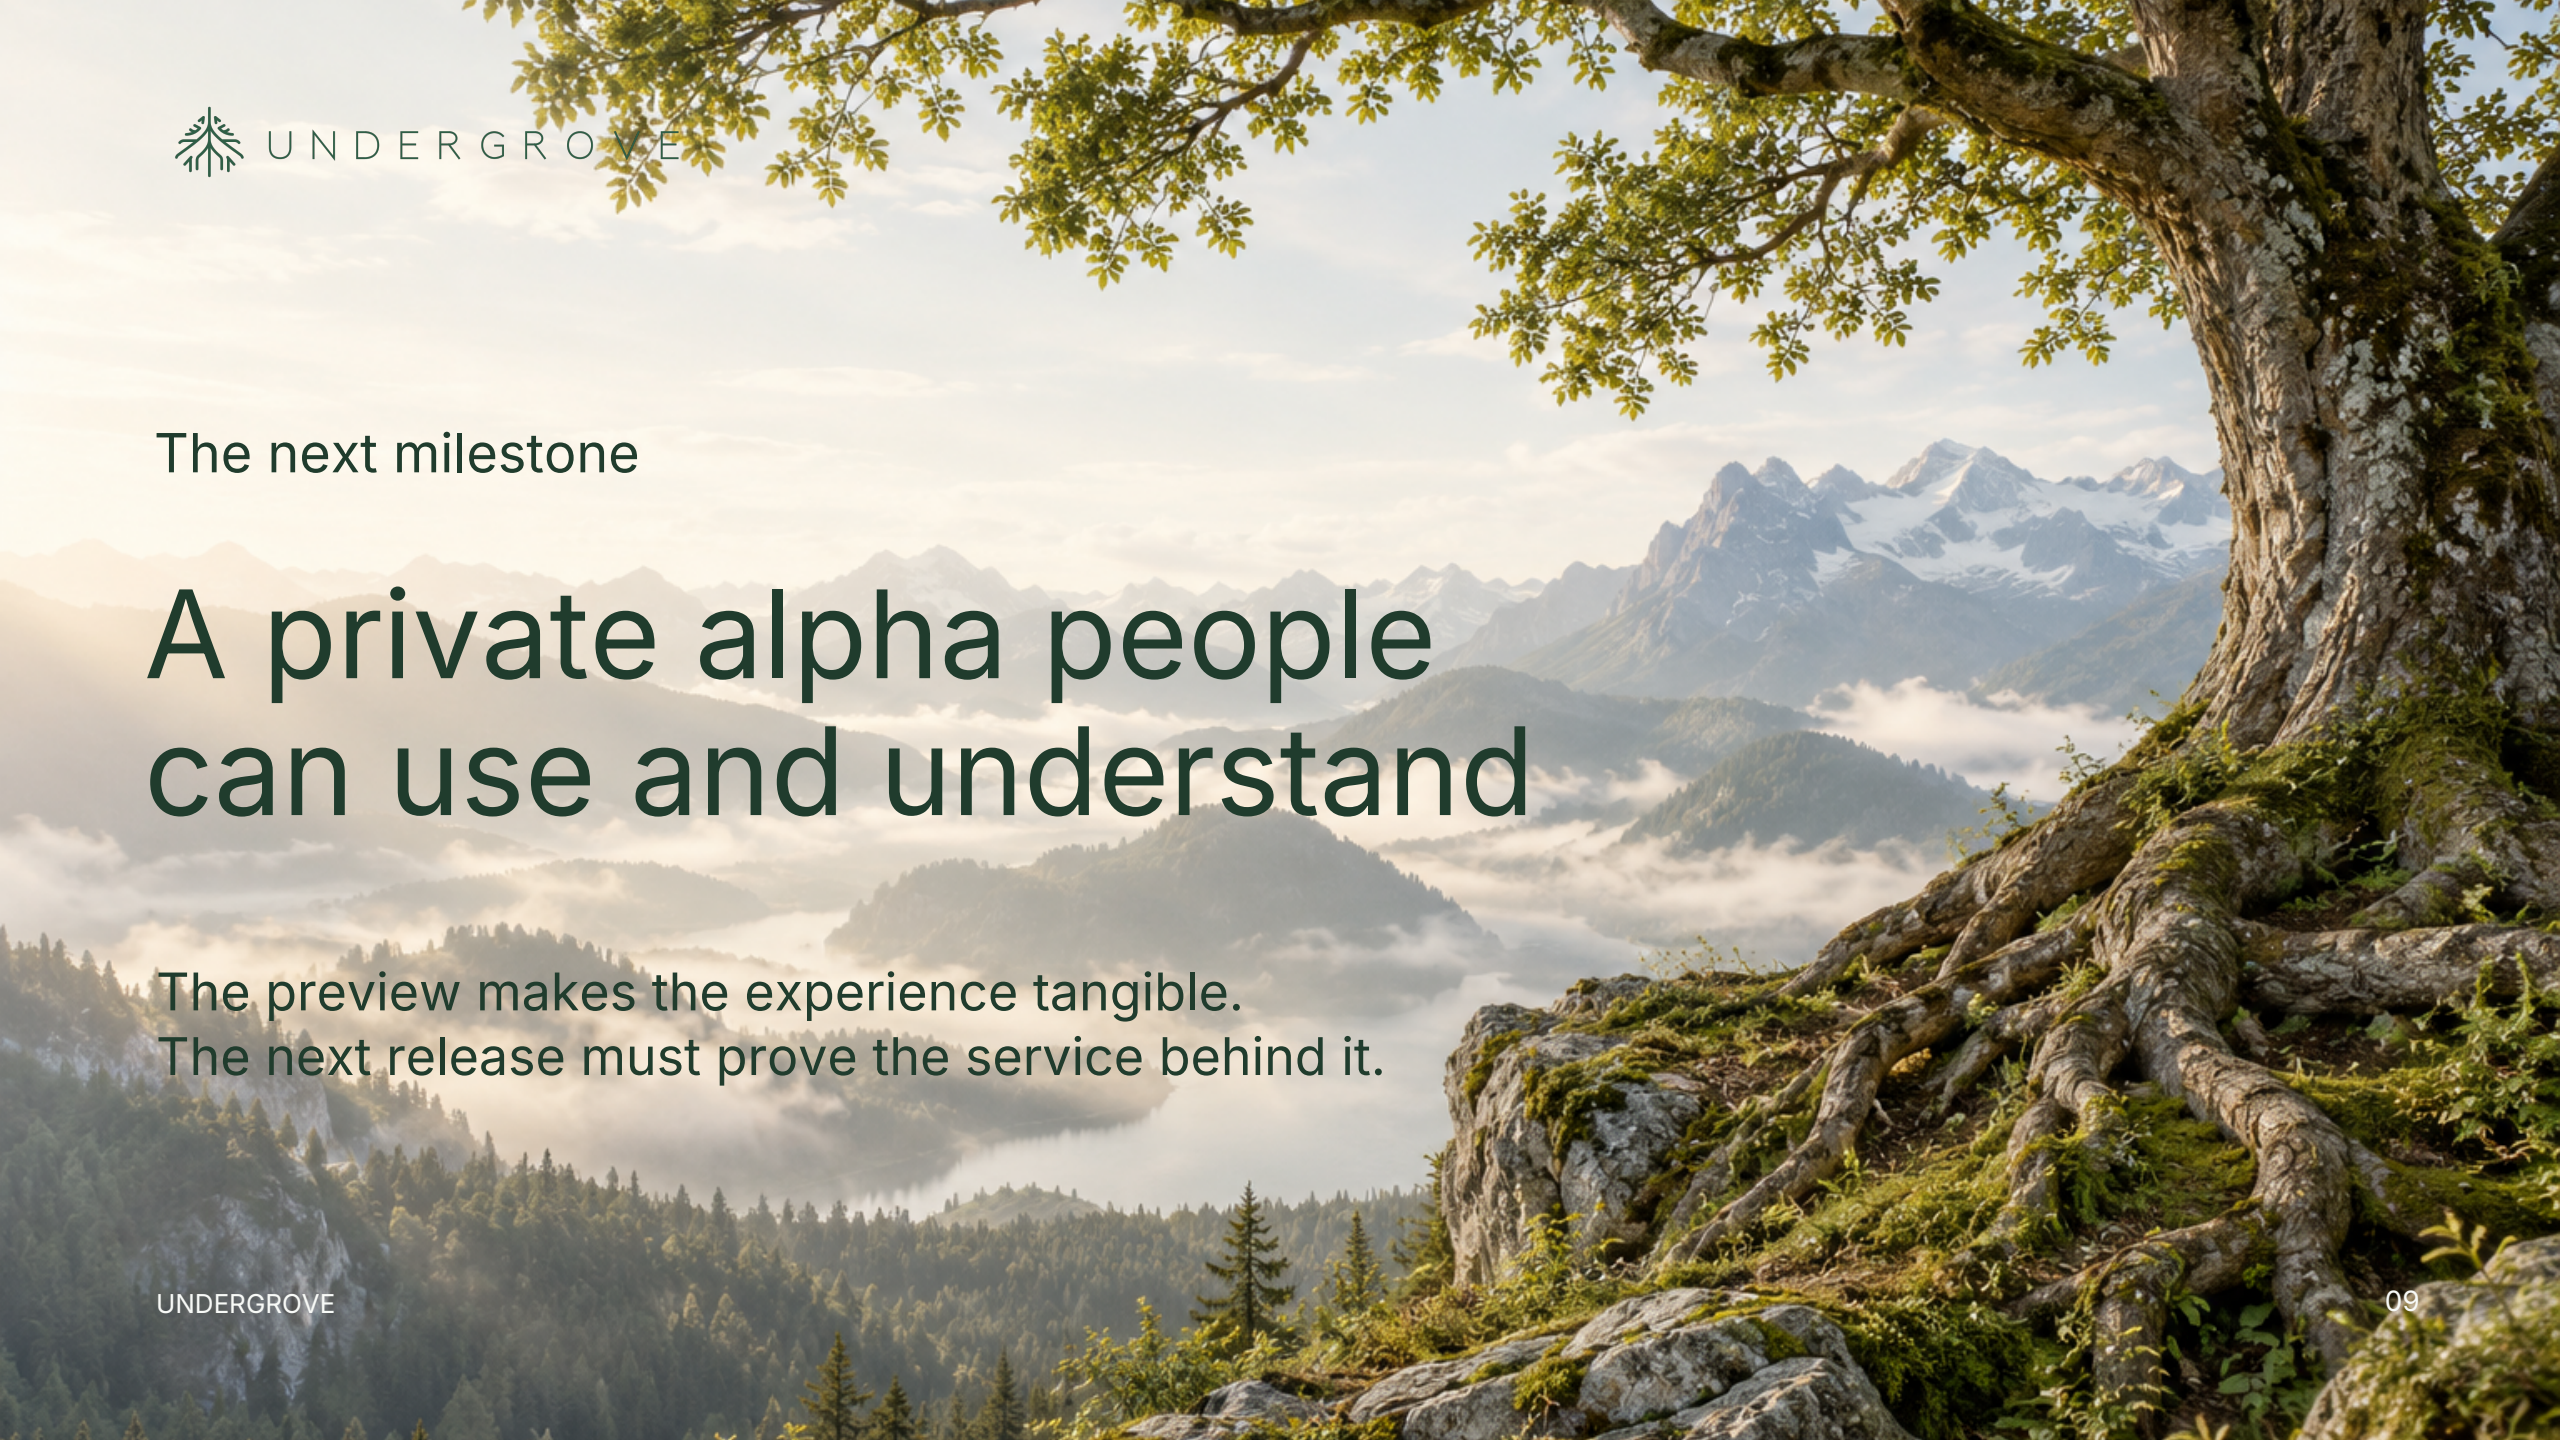
01	02	03
Private alpha	**Personal workspace**	**Paid beta**
One model and one hosting route	Context that users control	A service people can rely on
Measured quality and latency. A tested retention policy.	Permission-scoped retrieval. Reliable export and deletion.	Real task costs and capacity. Account recovery and support.

Additional model routes, tools and confidential computing follow separate reviews.

What the alpha needs to prove

Usefulness	Can people complete meaningful tasks and return?
Understanding	Can they explain where a chosen model sends data?
Economics	What does a completed task cost to serve?
Reliability	Do failures, access checks and deletion behave correctly?

Pricing follows observed usage and measured cost. Demand still needs validation.



The next milestone

A private alpha people can use and understand

The preview makes the experience tangible.
The next release must prove the service behind it.