Dweve
Weaving Intelligence

ENGLISH EDITION · JULY 2026

MARKETING BROCHURE · 16 PAGES

Research with a path to product

Explore ambitious mechanisms while keeping evidence, status, and product claims distinct.

VISUAL BRIEFING

RESEARCHERS, ENGINEERING PARTNERS, FUNDERS, AND PRODUCT REVIEWERS

ENGLISH

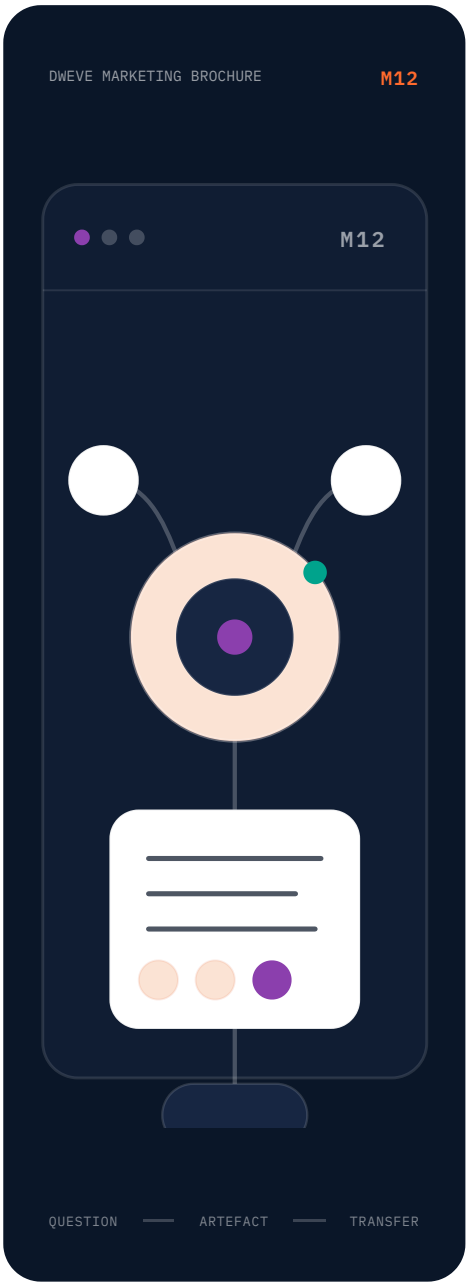

A collaboration brief for research, validation, and responsible transfer.

M12

THE PROPOSITION

A question that can survive scrutiny

AI is everywhere now, but most people do not know how it works, where their data goes, or whether they can trust the answers. Our research exists to fix that. We work on three problems that matter to everyone who uses technology.

How to read this visual: follow a falsifiable question through the proposition.

MECHANISM

Read the promise against the operating reality

AI is everywhere now, but most people do not know how it works, where their data goes, or whether they can trust the answers. Our research exists to fix that. We work on three problems that matter to everyone who uses technology.

Before research becomes a feature, we publish it so others can check our method, our data, and our conclusions. Open scrutiny catches mistakes and makes the work stronger.

- Question
- Mechanism
- Artefact
- Evidence

QUESTION TO CARRY

Which part of this proposition changes the outcome people receive today?

USE AND FIT

Where this proposition can earn its place

Start with responsibilities that have a real owner, useful sources, and a consequence worth improving.

How to read this visual: follow a falsifiable question through use and fit.

DECISION

Separate useful scope from attractive distraction

A confident answer is easy to give. A useful answer shows why it deserves your trust. We compare the sources, point out what they do not settle, and update the answer when the facts move. You can use the conclusion without losing the route back to its reasons.

The cost of a weak decision case often appears after approval. Teams rebuild the evidence, remember assumptions differently, and discover dependencies only when the operation changes. That recovery work is part of the original exposure, even when it is absent from the business case.

- Question
- Mechanism
- Artefact
- Evidence

QUESTION TO CARRY

Where does this earn a place, and where should it deliberately not be used?

THE PATH

From first question to an informed decision

A focused engagement should make the next decision easier, not create a larger promise.

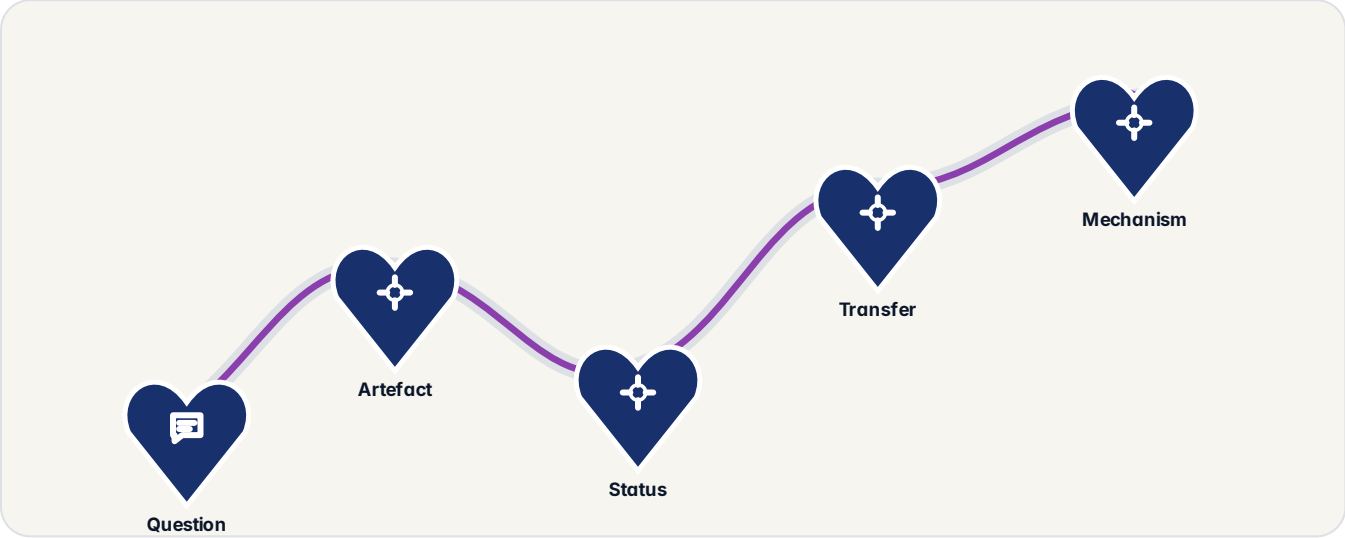

How to read this visual: follow a falsifiable question through the path.

WHY THIS SHAPE WORKS

The workflow, operating boundary, evidence, and decision criteria are considered together from the start.

MECHANISM

Make every stage earn the next one

Start with one decision where risk, cost, and operational impact are hard to separate. We can help make the question, evidence, assumptions, and change triggers explicit.

The first job is not to sound certain. It is to understand what you need to decide, which part of the question matters now, and what could make today's answer wrong tomorrow. That boundary keeps a simple explanation honest.

- Question
- Artefact
- Status
- Transfer

QUESTION TO CARRY

What evidence must exist before the scope is allowed to grow?

THE VALUE PICTURE

One outcome, seen from three responsibilities

Good AI work must make sense to the person using it, the team operating it, and the people reviewing it.

How to read this visual: follow a falsifiable question through the value picture.

PRIMARY READER

Researchers, engineering partners, funders, and product reviewers

The first conversation.

DECISION SUPPORTED

A collaboration brief for research, validation, and responsible transfer.

The next useful step.

DECISION

Read the same outcome from several responsibilities

A correction should not overwrite the past or leave downstream work stranded. A new source event identifies affected claims and outputs, reopens only the relevant review, and records the transition from earlier version to current version.

The goal is not a larger report. It is a tighter operating loop: frame the uncertainty, gather evidence, challenge the claim, make a bounded choice, observe what happened, and revise when reality disagrees.

- Question
- Artefact
- Status
- Mechanism

QUESTION TO CARRY

Can the same result satisfy its user, operator, and reviewer?

THE OPERATING MODEL

Control follows the work, from outcome to exit

A credible proposition connects the visible result to the less visible responsibilities underneath it.

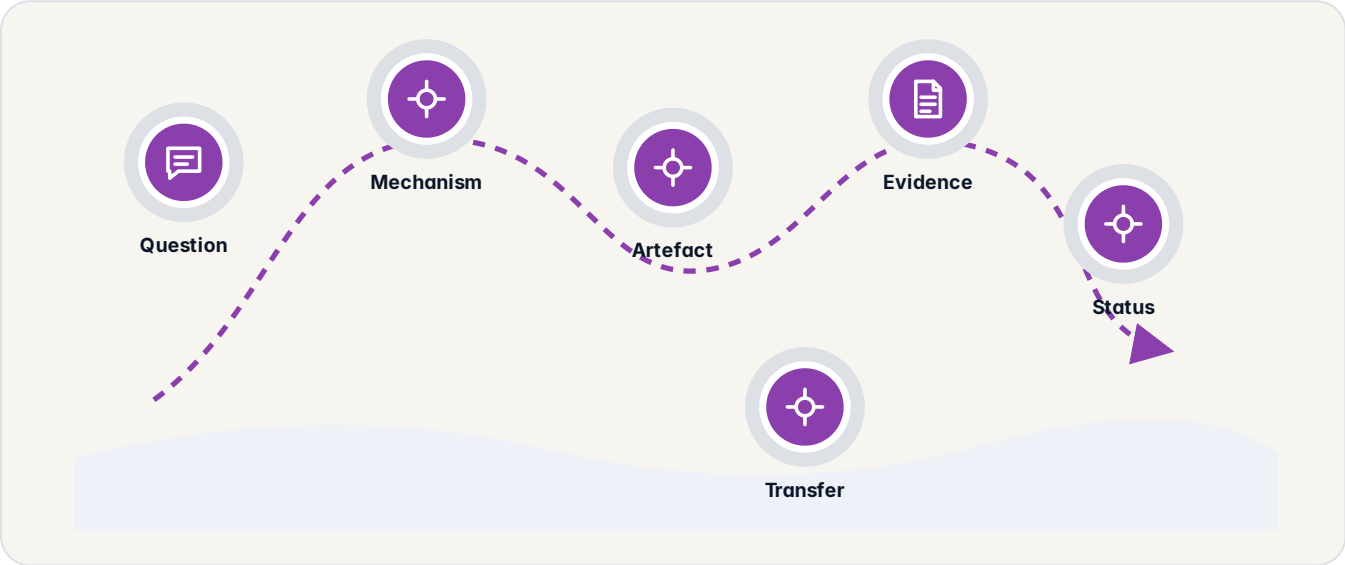

How to read this visual: follow a falsifiable question through the operating model.

MECHANISM

Name the controls behind the simple interface

Research into quantisation and hardware-specific kernel work that cuts the power a model needs without degrading the output. Less electricity means lower running cost and a smaller footprint.

Every material claim should reach a visible gate before the choice hardens. The gate tests direct support, a serious contrary view, and the assumptions the operation must keep true. A hold is useful when it names exactly what must be resolved.

- Question
- Mechanism
- Artefact
- Evidence

QUESTION TO CARRY

Who owns each control when the interface appears simple?

THE RESPONSIBILITY FLOW

People, platform, and evidence move together

The work becomes easier to govern when each stage leaves a clear hand-off.

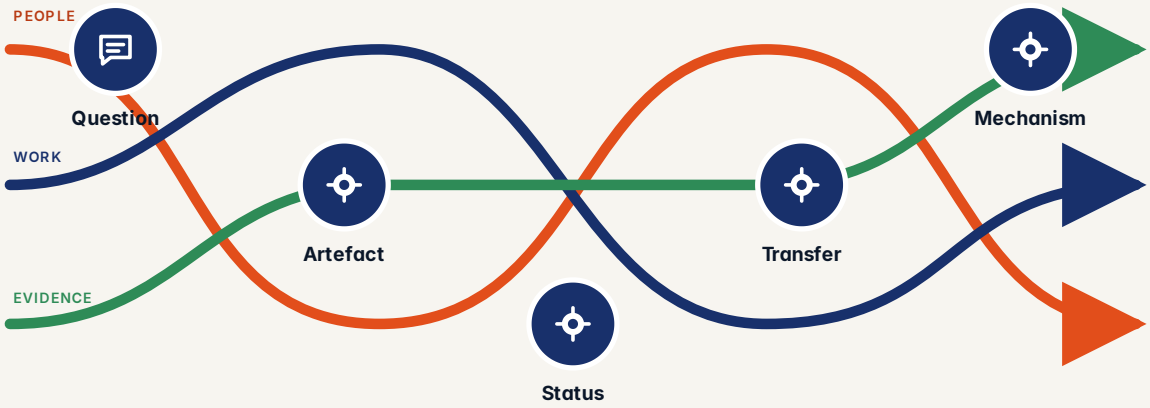

How to read this visual: follow a falsifiable question through the responsibility flow.

DECISION

Keep every hand-off attached to an owner

You should not need to read a methods paper before asking a normal question. The work shows up as a source you can open, a clear admission when the evidence is thin, and a visible update when the situation changes.

Decision confidence comes from knowing what supports the choice, which objection still matters, who owns the consequence, and what should trigger review. We build that case while the decision can still move, then keep the evidence attached after approval.

- Question
- Artefact
- Status
- Transfer

QUESTION TO CARRY

Is every hand-off visible, bounded, and assigned to an owner?

SOURCE TO RESULT

Every useful outcome begins with an authoritative basis

The work stays explainable when source identity, transformation, decision, and hand-off remain visible between systems.

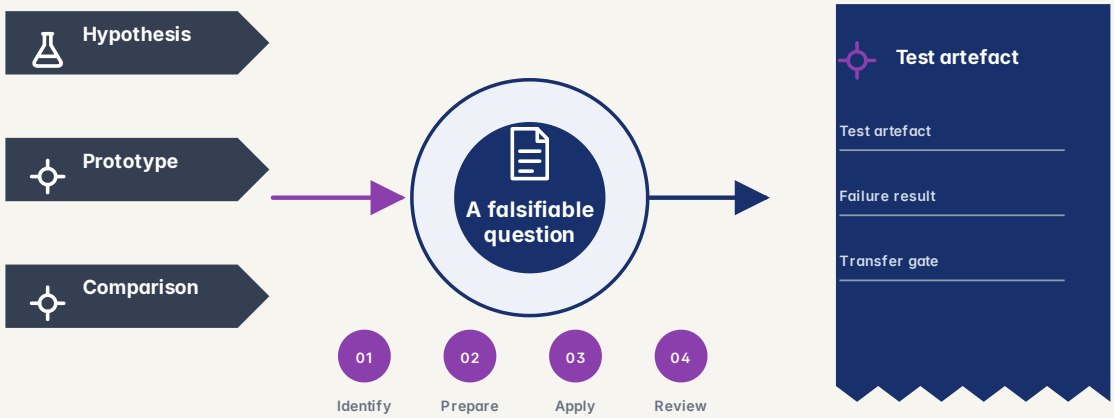

How to read this visual: follow a falsifiable question through source to result.

MECHANISM

Keep the basis visible as the work changes form

A result inherits every source and transformation behind it. Source identity, collection context, procedure version, and downstream dependencies enter the graph at the moment they are known so lineage remains queryable in both directions.

Several sources can repeat the same origin. A newer source can narrow an older conclusion. A serious objection can matter even when it is outnumbered. Research keeps those relationships visible so the decision owner can see where confidence depends on a fragile consensus.

- Identify
- Prepare
- Apply
- Review

QUESTION TO CARRY

Can the result still be connected to the exact basis that shaped it?

EVIDENCE PACKET

The record should answer the next difficult question

Capture enough to explain the result, inspect authority, investigate failure, and make the next decision without rebuilding the story from memory.

How to read this visual: follow a falsifiable question through evidence packet.

DECISION

Preserve the record needed by the next question

A conclusion should carry enough structure for another reader to inspect what was asked, which method ran, which inputs mattered, what was observed, where the limitations sit, and which lineage connects the parts. The packet is evidence of the process, not proof that the conclusion can never be wrong.

A source can help without settling the whole question. We separate the record closest to the event from background context and from evidence that points another way. The result explains which part is supported and which part is still open.

- Purpose
- Basis
- Authority
- Action

QUESTION TO CARRY

Does the retained record answer the next difficult question without reconstruction?

DEPLOYMENT CHOICE

Choose the operating posture around the work

Location is only one dimension. Decide who operates, who holds keys, which dependencies exist, how updates arrive, and what continuity requires.

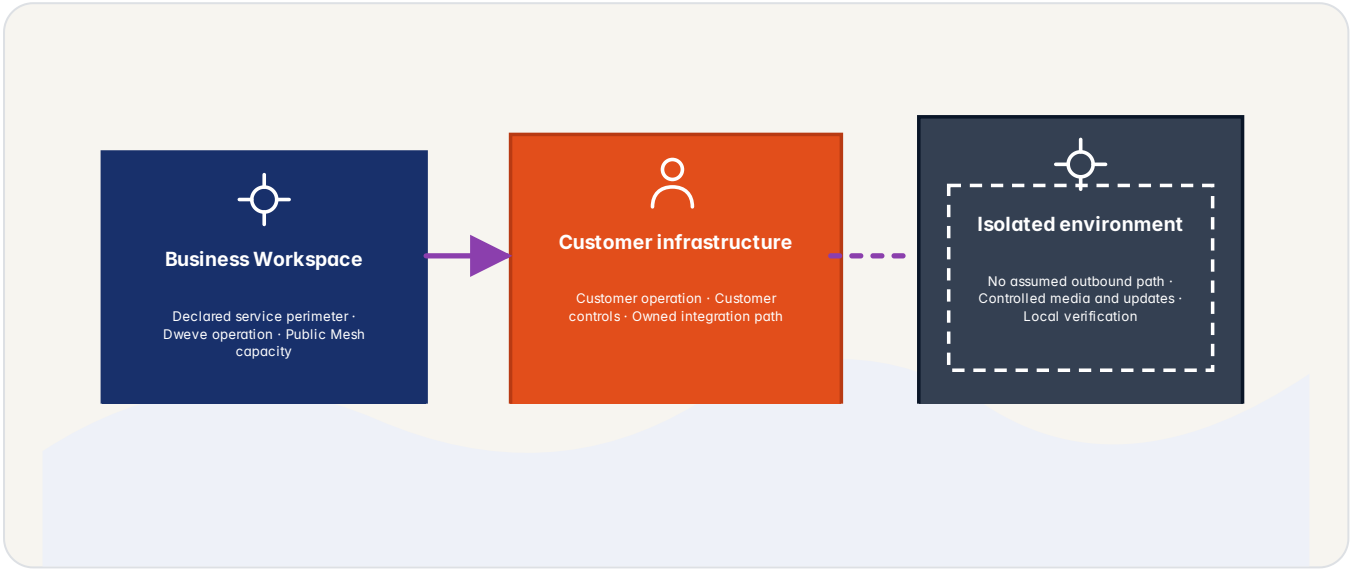

How to read this visual: follow a falsifiable question through deployment choice.

MECHANISM

Compare operating responsibility, not location alone

Transfer asks what the capability changes, what it costs to own, who operates it, and how failure is contained. Research value becomes business value only after those obligations are visible.

We study how to make AI explain its reasoning, give the same answer every time, and say so when it is unsure. You should never have to trust a black box. You should be able to see why it said something and decide for yourself.

- Begin through managed Fabric
- Run on customer infrastructure
- Close the environment
- Question

QUESTION TO CARRY

Who holds infrastructure, keys, updates, logs, recovery, and the exit path?

WHAT TO LOOK FOR

Three signals that the proof is earning trust

Progress is a useful outcome under understood control, supported by reviewable evidence.

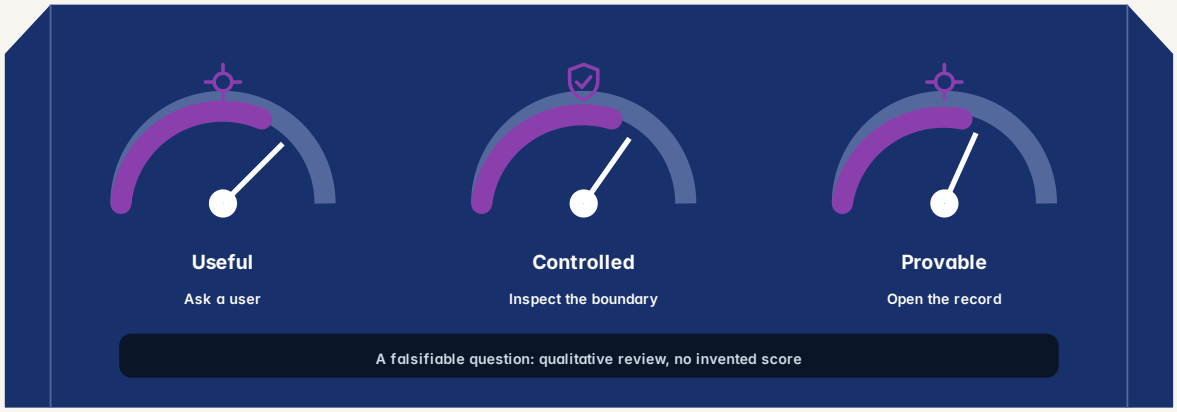

How to read this visual: follow a falsifiable question through what to look for.

DECISION

Use signals that can change the next decision

The tools you meet through Fabric should not ask you to trust a hidden leap. Our research explores how systems reason, adapt, compress, and replay so useful behaviour can arrive with an explanation close by.

Reproducibility is the foundation of every claim we make. Each row pairs a result with the stream that produced it and the figure we measured. Read it as a spec sheet, not a brochure.

- Useful
- Controlled
- Provable
- Question

QUESTION TO CARRY

Which signals would change the next decision rather than merely impress the room?

HONEST LIMITS

Design failure before the first success demo

Visible failure is more useful than a polished result whose limits remain unknown.

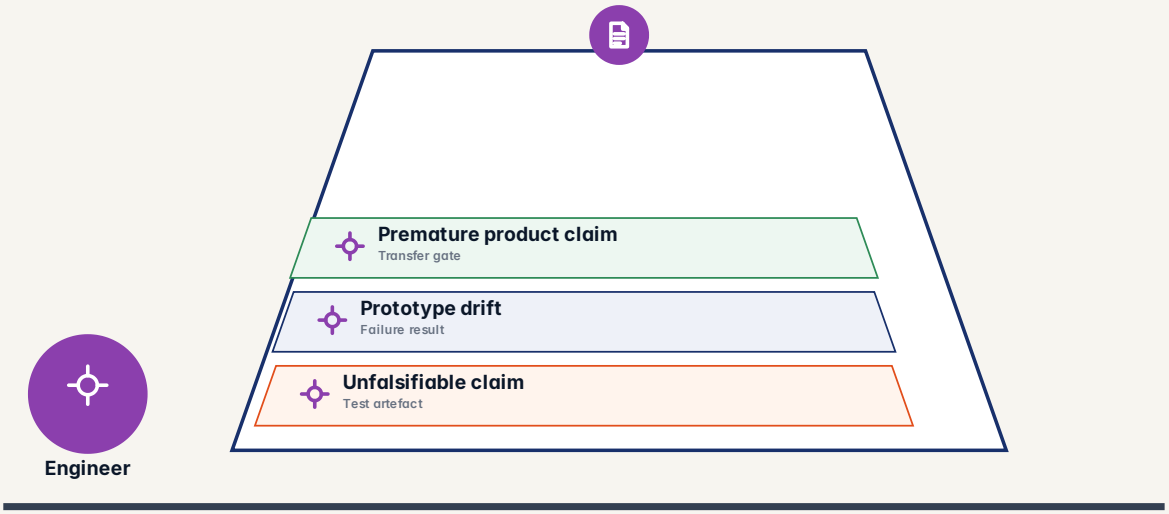

How to read this visual: follow a falsifiable question through honest limits.

MECHANISM

Design the failure before presenting the success

A promising result does not automatically become a feature. It crosses only when the benefit is clear, the behaviour remains understandable, and the failure mode can be handled in the product.

The method makes spend, evidence, challenge, and handoff visible. A programme can be continued, narrowed, or stopped without pretending that every interesting result deserves a product roadmap.

- Unfalsifiable claim
- Prototype drift
- Premature product claim

QUESTION TO CARRY

What happens when a source, permission, target, model, or supplier is unavailable?

ACCEPTANCE PATH

A proof should close with a decision, not applause

Agree the decision path before the demonstration so a strong or weak result can change what happens next.

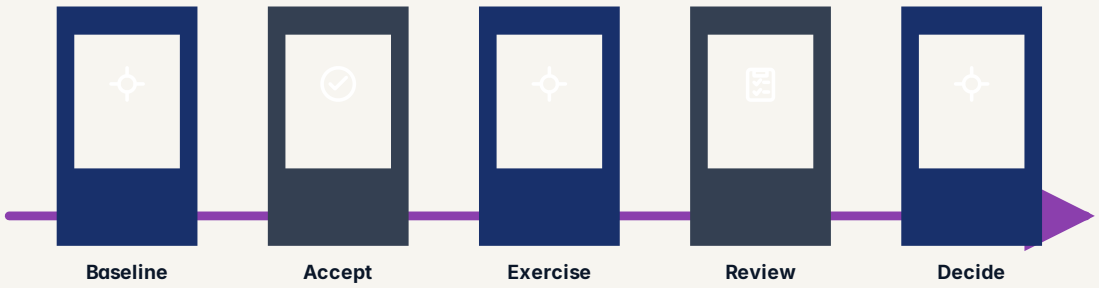

PROCEED · NARROW · HOLD · REJECT

How to read this visual: follow a falsifiable question through acceptance path.

DECISION

Close the proof with an owned decision

Declare the unit of enquiry, expected observation, acceptance rule, and kill condition before a run can produce an answer. The contract prevents the hypothesis from moving after the result and gives a replay an explicit target.

We name what someone needs to understand, what must remain fixed, and what would change the answer. That boundary keeps a demonstration from quietly becoming a promise.

- Baseline
- Accept
- Exercise
- Review

QUESTION TO CARRY

Are success, counter-metrics, failure cases, and stop conditions agreed first?

ADOPTION

Expansion is a sequence of owned choices

Move from one bounded result to a repeatable operating capability without making the scope vague.

How to read this visual: follow a falsifiable question through adoption.

MECHANISM

Expand only what proved reusable

Dweve Research separates promising mechanisms from expensive assumptions. Each programme starts with a bounded operational question, records what survives testing, and shows what would have to be true before a result can move into practice.

Programme pages expose the mechanism and research website description. Whitepapers carry the longer arguments. Use both to separate an interesting capability from one ready to enter an operating plan.

- Prove one workflow
- Own the operation
- Repeat the pattern
- Add responsibility

QUESTION TO CARRY

What can be reused safely, and what requires a new decision and proof?

WHAT TO PREPARE

Bring the ingredients for a useful first session

A small amount of real context produces a better conversation than a broad catalogue of hypothetical features.

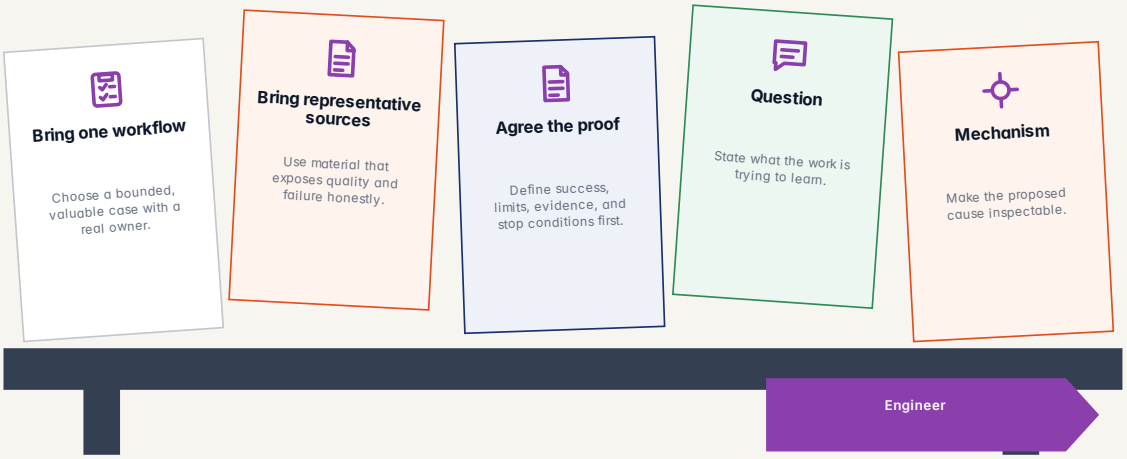

How to read this visual: follow a falsifiable question through what to prepare.

KEEP IT BOUNDED

Sensitive production data is not required for the first conversation. Bring representative material and the rules that define a useful result.

DECISION

Bring enough reality to make the session useful

Dweve is built by a small founding team in Arnhem that pairs deep industry engineering with data governance and decades of teaching. The research is grounded in real production constraints, and the products are grounded in verified results.

We build prototypes, run experiments, and measure the results against real data. If something does not work, we write that down too. A negative result is still a result.

- Bring one workflow
- Bring representative sources
- Agree the proof
- Question

QUESTION TO CARRY

Which real sources, rules, owners, and examples will make the first session honest?

THE INVITATION

Continue with the question that matters

Scope a briefing, demonstration, or proof around the workflow and decision that matter to your organisation.

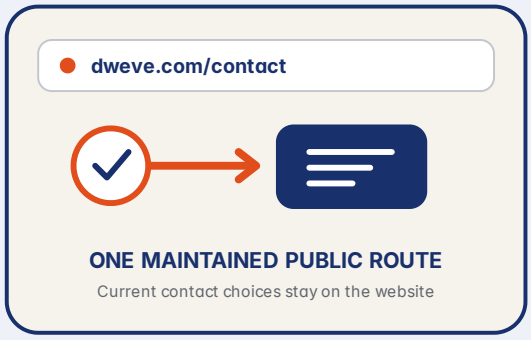

CONTINUE ONLINE

dweve.com/contact

Current route and contact choices

START A CONVERSATION

Continue through the current contact page

Scope a workflow or ask a technical question, then use the contact page to begin.

START	**dweve.com/contact**	Use the maintained public contact route
BRING	**One bounded question**	Workflow, audience, decision, and desired evidence
CHOOSE	**Briefing · demonstration · PoV**	Select the smallest useful next step
VERIFY	**Current scope and terms**	Confirm availability, pricing, and responsibilities before commitment

MECHANISM

Turn the brochure into one bounded next step

Reproducibility starts before execution. We model the question as a falsifiable contract, preserve provenance at collection, identify every transformation and runtime condition, and connect observations to downstream claims. A replay should either reproduce the path or expose the condition that moved.

The research portfolio does not force every question through one method. Jacquard studies deterministic, typed worlds. Forge studies program synthesis for kernels with solver backed verification. Mycelia studies artificial life and emergence in a deterministic simulation.

- Start: dweve.com/contact
- Bring: One bounded question
- Choose: Briefing · demonstration · PoV
- Verify: Current scope and terms

QUESTION TO CARRY

What is the smallest useful next conversation Dweve should prepare?