

MARKETING BROCHURE · 16 PAGES

Fabric: the product brochure

A standalone visual introduction to Fabric, the job it owns, how it works, where it fits, and how to prove it.

VISUAL BRIEFING

PRODUCT LEADERS, BUYERS, ARCHITECTS, DELIVERY TEAMS, AND TECHNICAL EVALUATORS

ENGLISH




Decide whether Fabric belongs in the selected workflow.

PR01

THE PROPOSITION

The workspace where AI becomes real work

Fabric is the workspace for questions that grow into real work. Ask, add a document, invite another kind of help, and keep the useful answer connected to its sources. When you share it, the evidence can travel with the result.



How to read this visual: follow a governed workspace through the proposition.

MECHANISM

Read the promise against the operating reality

Most AI apps keep the conversation inside one assistant. Need another opinion? Open another app. Need a different kind of AI? Upload the file again. Need to compare the answers? Copy them into a third place. Fabric keeps the task together. You can ask a question, add a document, invite another kind of AI, bring in a helper, save the useful process, and share the result with its source attached. The helper can change. The work does not start over.

Fabric is the workspace for questions that grow into real work. Ask, add a document, invite another kind of help, and keep the useful answer connected to its sources. When you share it, the evidence can travel with the result.

- Ask
- Bring
- Compose
- Govern


QUESTION TO CARRY

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

USE AND FIT

Where Fabric 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 governed workspace through use and fit.

DECISION

Separate useful scope from attractive distraction

Once your access group opens, start with one real task and add a document or another kind of help when the work needs it. The conversation stays connected to the sources, decisions, and reusable steps that follow.

A useful conversation may reveal a process worth keeping. Fabric flows make the process explicit. A trigger can begin in chat, through a schedule, from a webhook, or through the API. Parallel branches can ingest and retrieve at the same time. Different nodes can use different models or deterministic components. Conditions can route the work.

- Ask
- Bring
- Compose
- Govern


QUESTION TO CARRY

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

THE PATH

How Fabric moves from question to decision

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



How to read this visual: follow a governed workspace 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

The approval packet is created as the work moves, not rebuilt later when someone asks. Source, rule, owner, and a replay path attach while the decision is made, so a regulator question is a retrieval, not a reconstruction.

Fabric is the workspace where AI work stops being scattered pilots and becomes one thing you can govern. On the first morning a team logs in to folders that already exist, approvals already routed, and an answer that arrives with its source and its sign-off attached. One dashboard for who can do what, what each team is spending, and a trail of every decision that survives a regulator question.

- Ask
- Compose
- Review
- Reuse


QUESTION TO CARRY

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

THE VALUE PICTURE

Fabric 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 governed workspace through the value picture.

PRIMARY READER

Product leaders, buyers, architects, delivery teams, and technical evaluators

The first conversation.

DECISION SUPPORTED

Decide whether Fabric belongs in the selected workflow.

The next useful step.

DECISION

Read the same outcome from several responsibilities

Fabric group conversations can contain distinct people, models, agents, and teams. Every contribution retains its author. Context can be shared with the whole room or projected only to the participant that needs it. A person can ask one participant directly, invite another model for a second view, request a verifier, or bring a domain team into the same thread.

Provider SDKs want to own the conversation, message format, tools, memory, streaming protocol, agent schema, and execution history. That is convenient for one integration and expensive for a platform. Fabric defines the durable system above those providers. The workspace owns the object graph. The work owns the context. Models, agents, teams, solvers, tools, and people participate through typed contracts. The UI, REST API, WebSocket stream, flow runtime, knowledge layer, model registry, and activity view operate on the same domain.

- Ask
- Compose
- Review
- Bring


QUESTION TO CARRY

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

THE OPERATING MODEL

The control model around Fabric

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



How to read this visual: follow a governed workspace through the operating model.

MECHANISM

Name the controls behind the simple interface

A human can override the automatic result with a recorded reason. A closed model can still be used as a judge, but Fabric records the fact that the judgement came from that model and does not upgrade its prose into formal proof. Best answer is never used without a selection basis.

A correction has an address. A policy change has a visible effect. A model change can be compared against the previous execution. Evidence is created as the operation runs rather than reconstructed after the dispute arrives. Fabric exposes the system decisions it controls and marks the provider boundaries it cannot see through.

- Ask
- Bring
- Compose
- Govern


QUESTION TO CARRY

Who owns each control when the interface appears simple?

THE RESPONSIBILITY FLOW

People, Fabric, 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 governed workspace through the responsibility flow.

DECISION

Keep every hand-off attached to an owner

A piece of work has identity separate from the model, agent, or person currently handling it. The same work can move from chat into a flow, pause at an approval, call a tool, invite a team, produce an output, and return as reusable knowledge without becoming unrelated records.

Roles, permissions, budgets, approvals, and policies attach to the resources and actions they govern. An owner can manage the workspace. An editor can work inside assigned scope. A viewer can read only what was shared. An auditor can query the evidence without receiving operational authority.

- Ask
- Compose
- Review
- Reuse


QUESTION TO CARRY

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

SOURCE TO RESULT

Fabric: from authoritative source to result

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



How to read this visual: follow a governed workspace through source to result.

MECHANISM

Keep the basis visible as the work changes form

The document can become part of your knowledge library. The useful steps can become a saved flow. A specialist helper can remain available for the next similar task. The answer can be shared with the source and relevant thread intact. Chat is one entrance into Fabric. It is not the whole product.

For each supported operation, Fabric can expose a composition record containing work and correlation identifiers, workspace and project scope, participant identities and roles, model IDs, versions, runtimes, and placement, knowledge nodes and version hashes, and retrieval and reranking data.

- Identify
- Prepare
- Apply
- Review

QUESTION TO CARRY

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

EVIDENCE PACKET

The evidence Fabric leaves with the work

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 governed workspace through evidence packet.

DECISION

Preserve the record needed by the next question

It can also represent training or fine-tuning support, region and residency constraints, cost and latency characteristics, determinism and replay properties, evidence and trace depth, evaluation and publication state, and permitted work classes. Selection can then operate on explicit capability and policy rather than a hard-coded dropdown. Model-agnostic does not mean pretending every model is equivalent.

A workflow can pin a model version. An agent can allow a set of eligible models. A project can prohibit external providers. A regulated flow can require local execution and a minimum evidence depth. A low-risk task can choose the smallest adequate model. The decision is recordable and appended to the work record.

- Purpose
- Basis
- Authority
- Action

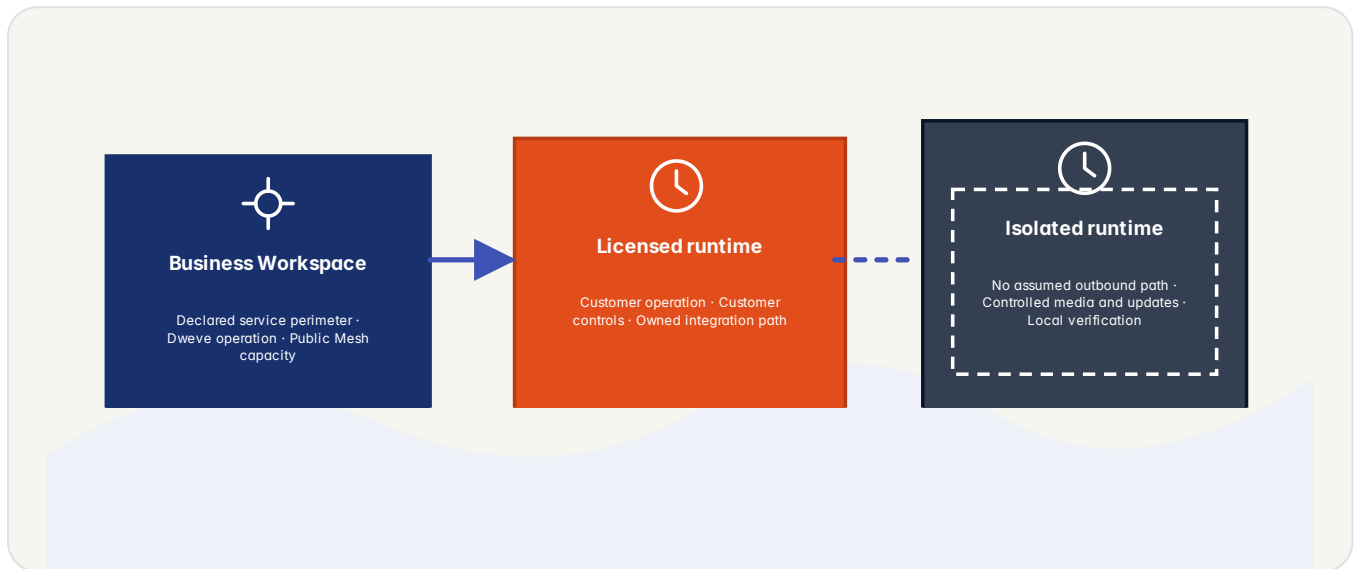
QUESTION TO CARRY

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

DEPLOYMENT CHOICE

Choose the operating posture for Fabric

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 governed workspace through deployment choice.

MECHANISM

Compare operating responsibility, not location alone

A model begins as one component. Then its chat format becomes the conversation layer, its memory becomes the organisation's memory, its tool API becomes the integration strategy, its agent format becomes the automation layer, its provider logs become the audit evidence, its pricing becomes the operating model, and its region becomes the residency posture. The model has acquired operational gravity.

Fabric is the control plane that turns AI from a set of experiments into a platform capability your board can steer. Talk to sales about a pilot, or scope a deployment on infrastructure you already run.

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


QUESTION TO CARRY

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

WHAT TO LOOK FOR

Signals that Fabric is earning trust

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



How to read this visual: follow a governed workspace through what to look for.

DECISION

Use signals that can change the next decision

Fabric records the differences rather than hiding them. Capabilities, context limits, tool support, region, cost, latency, evidence depth, determinism, local availability, training support, and policy eligibility can all affect where a model may be used. The organisation selects intelligence for the work instead of rebuilding the work for the provider.

A screenshot carries the conclusion but drops the evidence. Fabric can share the selected answer with the relevant source and thread attached. The next person can check what supported it without receiving the rest of your workspace.

- Useful
- Controlled
- Provable
- Ask

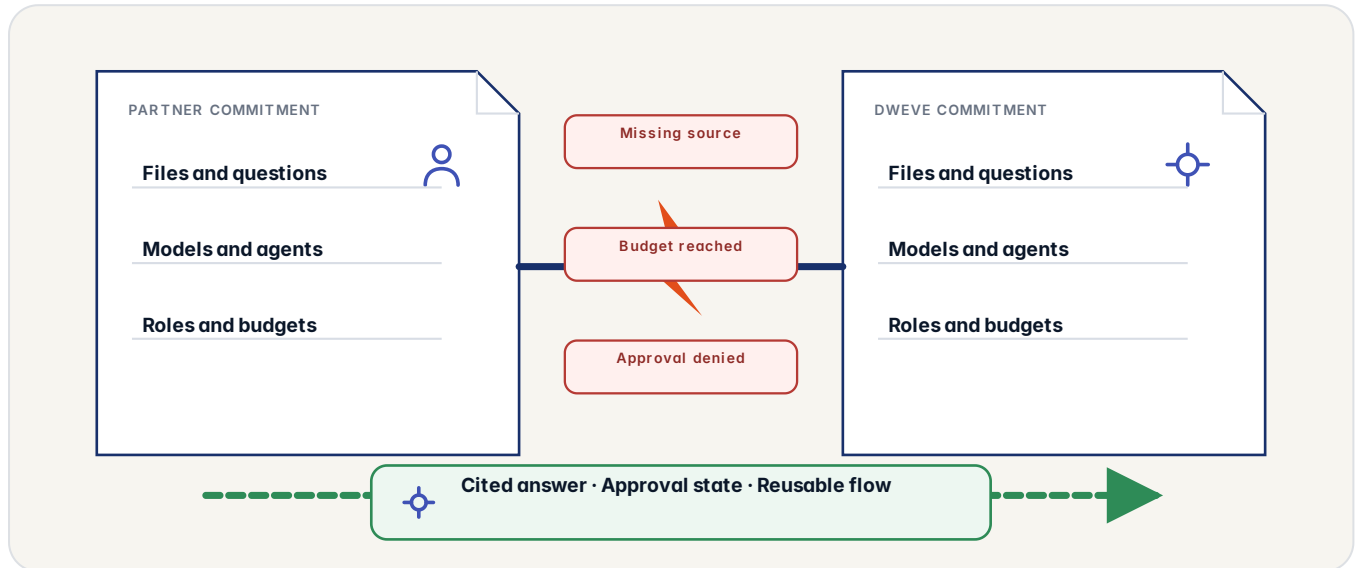
QUESTION TO CARRY

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

HONEST LIMITS

Test how Fabric fails 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 governed workspace through honest limits.

MECHANISM

Design the failure before presenting the success

Fabric brings five surfaces together under one contract: the agent roster, the workflow engine, the model layer, the knowledge base, and the governed workspace. Each surface has its own API namespace. Each surface writes to the same tenant audit chain. You use only the surfaces you need, and each one is available from the first tier.

Approve with one condition: deployment must stay inside the agreed EU boundary. The residency clause in v4 names a fallback region that fails the DORA third-party resilience test, so the condition pins it back. The clause, the policy version, and the reviewing manager are recorded, and the decision replays the same way for an audit.

- Missing source
- Budget reached
- Approval denied


QUESTION TO CARRY

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

ACCEPTANCE PATH

Close the Fabric proof with a decision

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



How to read this visual: follow a governed workspace through acceptance path.

DECISION

Close the proof with an owned decision

When Fabric selects or combines candidates, the basis should be inspectable. A judge can use explicit criteria. A weighting step can record its inputs and weights. A constraint stage can name the constraints evaluated. A solver can return a decision and, where supported, a certificate or unsat core.

Fabric can preserve the candidate set, the selection criteria, the weights, the judge, and the final decision. A model council is therefore not a row of answers on a comparison screen. It is a governed collaboration inside the work, where every contribution remains attributable in the record.

- Baseline
- Accept
- Exercise
- Review


QUESTION TO CARRY

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

ADOPTION

Adopt Fabric through owned choices

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



How to read this visual: follow a governed workspace through adoption.

MECHANISM

Expand only what proved reusable

Fabric treats your code as a first-class citizen of the platform. Register custom workflow nodes and custom tools so in-house systems appear beside the built-in ones. Your engineers extend Fabric the same way the platform ships, typed, versioned, and observable. A custom node behaves like a platform node in the visual editor, and a custom tool is offered to agents with typed parameters and structured returns.

A Dweve capability may be better when you want inspectable reasoning. A saved helper may know your usual process. A person may need to approve the final step. Fabric lets them contribute without splitting the task into unrelated chats. The files, questions, answers, sources, and decisions remain together.

- 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

What to bring to the Fabric working 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 governed workspace 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

Upload your files, ask a real question, and get an answer tied back to the exact source. Fabric is strongest when the answer stays grounded and you can check it.

A Fabric group conversation is a protocol surface, not a role-play prompt. Participants have stable identities, capabilities, permissions, context scopes, and message attribution. A participant may be a person, model, agent, team, judge, verifier, retriever, or tool-backed worker.

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

QUESTION TO CARRY

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

THE INVITATION

Continue the Fabric conversation

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

●
dweve.com/contact

ONE MAINTAINED PUBLIC ROUTE

Current contact choices stay on the website

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

Replacing it later is no longer a model decision. It is a platform migration. Fabric separates the durable AI system from the models currently serving it. The organisation keeps its knowledge, agents, teams, workflows, integrations, roles, history, and evidence while the model portfolio evolves.

Start with a piece of work that currently crosses chats, documents, AI providers, agents, spreadsheets, approvals, and systems. See how Fabric keeps its context, participants, sources, decisions, and actions together while the intelligence underneath remains replaceable.

- 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?