

MARKETING BROCHURE · 16 PAGES


Loom: the product brochure

A standalone visual introduction to Loom, 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 Loom belongs in the selected workflow.

PR02

THE PROPOSITION

Cognition composed for the task

A model can predict. The task may also need perception, memory, retrieval, ranking, reasoning, a solver, or a verifier. Loom activates the relevant cognitive threads, connects their typed results, and records the path that produced the answer. A model is one thread. Loom is the weave.



How to read this visual: follow a composed cognition path through the proposition.

DECISION

Read the promise against the operating reality

A model can predict. The task may also need perception, memory, retrieval, ranking, reasoning, a solver, or a verifier. Loom activates the relevant cognitive threads, connects their typed results, and records the path that produced the answer. A model is one thread. Loom is the weave.

A new capability should not require retraining the entire mind. Loom can add a classifier, a solver, a reasoning mode, a memory procedure, a verifier, or a domain expert as a new thread with its own tests and version. Existing threads do not need to be rewritten merely because the catalogue expanded. A new solver can handle a new problem class. A new expert can add domain depth. A new perception stage can make another modality available. A new micro model can replace a broad generative path for a narrow task.

- Perceive
- Remember
- Retrieve
- Reason


QUESTION TO CARRY

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

USE AND FIT

Where Loom 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 composed cognition path through use and fit.

MECHANISM

Separate useful scope from attractive distraction

Do not route every request to one model. Represent perception, memory, models, reasoning, solvers, verification, and experts as the distinct computational responsibilities they are. Then let Loom weave only the graph the task needs.

You should never have to wonder whether Loom had a bad day. Ask the same question with the same settings and you get the same answer, word for word. The specialists show their work so you can check before you trust. And because inference is fast and runs on the hardware you already have, you stop rationing questions. Ask freely. Experiment. Loom keeps up.

- Perceive
- Remember
- Retrieve
- Reason


QUESTION TO CARRY

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

THE PATH

How Loom 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 composed cognition path through the path.

WHY THIS SHAPE WORKS

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

DECISION

Make every stage earn the next one

Loom is not a model router with a larger catalogue. It is a runtime for composing heterogeneous cognitive components into one typed execution graph. A node can be a perception transform, a memory operation, a micro model, an embedding or retrieval stage, a reasoning mode, a constraint set, a native or integrated solver, a knowledge operation, a proof checker, a domain expert, or an expression stage. The planner resolves the graph from the task, policy, available capability, evidence requirements, and execution budget. Only the required graph becomes active.

A class decision is hidden in prose. An extracted clause is hidden in prose. A relevance judgement is hidden in prose. A calculation is hidden in prose. A constraint check is hidden in prose. A proof obligation is hidden in prose. A refusal is hidden in prose. Loom restores the boundaries. Each cognitive thread has a defined responsibility, input, output, version, and evaluation surface. The answer can still be simple. The system behind it no longer has to be vague.

- Perceive
- Retrieve
- Verify
- Express


QUESTION TO CARRY

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

THE VALUE PICTURE

Loom 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 composed cognition path through the value picture.

PRIMARY READER

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

The first conversation.

DECISION SUPPORTED

Decide whether Loom belongs in the selected workflow.

The next useful step.

MECHANISM

Read the same outcome from several responsibilities

A system that can weave many capabilities must also know when none of them can support a defensible result. Loom can surface that state explicitly. It does not need to force every cognitive path into a fluent answer. The perception may be incomplete. The memory may not contain the required fact. Retrieval may return weak evidence. The constraints may conflict. The solver may return unknown. The experts may disagree.

Loom can work with more than typed text. It can inspect documents, tables, images, audio, video, and structured information where the product supports them. It can notice layout, find names and dates, recognise patterns, inspect fields, and turn the input into something the next part of the system can use.

- Perceive
- Retrieve
- Verify
- Remember


QUESTION TO CARRY

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

THE OPERATING MODEL

The control model around Loom

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



How to read this visual: follow a composed cognition path through the operating model.

DECISION

Name the controls behind the simple interface

A transformer can write a plausible answer to a scheduling problem. That does not mean the schedule satisfies every constraint. A language model can guess whether a formula is satisfiable. That is not a SAT result. A model can describe an optimisation trade-off. That is not the optimum. Loom treats solvers as first-class cognitive threads.

A general-purpose model gives you a plausible answer and no way to defend it. Loom has 528 top-level domain experts and selects around 4 to 8 on an average request. The actual weave follows the task, and every output records the specialists and other cognitive threads that were active.

- Perceive
- Remember
- Retrieve
- Reason


QUESTION TO CARRY

Who owns each control when the interface appears simple?

THE RESPONSIBILITY FLOW

People, Loom, 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 composed cognition path through the responsibility flow.

MECHANISM

Keep every hand-off attached to an owner

Search contains several separate judgements. The query must be represented. Candidates must be found. Relevance must be estimated. Results must be ranked. The strongest candidates may need reranking. Policy, recency, provenance, or domain constraints may need to modify the final order. Loom gives each stage its own thread.

Classification, exact span extraction, relevance, ranking, reranking, retrieval, and embeddings are examples, not the boundary of the model family. Loom can host purpose-built learned components wherever a bounded statistical decision is the correct abstraction. Each model owns one responsibility and one evaluation surface.

- Perceive
- Retrieve
- Verify
- Express


QUESTION TO CARRY

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

SOURCE TO RESULT

Loom: 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 composed cognition path through source to result.

DECISION

Keep the basis visible as the work changes form

Some work requires more than a general cognitive path. Loom includes a catalogue of 528 experts across foundational reasoning, mathematics, science, code and systems, language, domain industries, multimodal work, verification, transfer, and metacognition. A query does not activate all of them.

A system cannot reason correctly about what it failed to perceive correctly. Loom includes perception systems for text, documents, images, audio, video, and structured data. Perception is not treated as a disposable prompt-preparation step. Its outputs remain addressable in the weave.

- Identify
- Prepare
- Apply
- Review


QUESTION TO CARRY

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

EVIDENCE PACKET

The evidence Loom 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 composed cognition path through evidence packet.

MECHANISM

Preserve the record needed by the next question

An intelligence system needs more than a message history. Loom separates memory by purpose. Retrieval is scoped to the task and the active cognitive thread. The memory system can consolidate, decay, forget, replay, and inspect its own state without turning memory into an unbounded transcript.

Loom does not run one fixed pipeline for all tasks. The task is classified by structure, modality, domain, required evidence, and possible solver family. The planner can construct a shallow graph for a bounded task or a deeper graph for a complex one. The graph is explicit enough to inspect and dynamic enough to fit the task.

- 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 Loom

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 composed cognition path through deployment choice.

DECISION

Compare operating responsibility, not location alone

The deployment determines where the weave runs. It does not redefine what the weave means. Available integrations and hardware change with the posture. The cognitive contract does not.

A Loom expert is trained on its domain, carrying its own knowledge graph, its own constraint set, and its own evaluation suite. Each expert carries 2 to 3.5 million constraints in its domain catalogue. Two experts working on the same token produce different outputs because they are different specialists, not because they were given different slices of the same neural network.

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


QUESTION TO CARRY

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

WHAT TO LOOK FOR

Signals that Loom is earning trust

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



How to read this visual: follow a composed cognition path through what to look for.

MECHANISM

Use signals that can change the next decision

A fluent answer is not automatically a supported answer. Loom can compare the result with the evidence, constraints, rules, checker results, and the checks used by the task. Where a formal certificate is available, the system can attach it.

Training, adaptation, routing, and control remain binary and integer. This common substrate makes the learned threads compact, deterministic on supported paths, and economical to activate. It is what allows Loom to carry many specialised capabilities without requiring all of them to run on every request. Symbolic and solver components retain their own native structures: graphs, intervals, proofs, and constraints are not forced into a neural representation.

- Useful
- Controlled
- Provable
- Perceive


QUESTION TO CARRY

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

HONEST LIMITS

Test how Loom 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 composed cognition path through honest limits.

DECISION

Design the failure before presenting the success

A span extractor emits offsets and labels. An embedder emits a binary representation. A retrieval stage emits candidates and distances. A ranker emits order and scores. A SAT solver emits SAT, UNSAT, or unknown, plus a model or certificate where supported. A causal reasoner emits paths, interventions, or independence results. A verifier emits an accepted result, a rejected result, or a named reason it cannot certify the path.

The solver fabric spans logical, constraint, optimisation, temporal, spatial, symbolic, graph, automata, equilibrium, probabilistic, and synthesis problems. When the problem has a solver-shaped answer, Loom can route to a solver. The result can be satisfiable, unsatisfiable, optimal, infeasible, proved, disproved, unknown, or bounded by the solver's actual contract. It does not need to become persuasive prose first.

- Missing context
- Failed check
- Unsupported expression


QUESTION TO CARRY

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

ACCEPTANCE PATH

Close the Loom proof with a decision

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 composed cognition path through acceptance path.

MECHANISM

Close the proof with an owned decision

The trace can name the perception outputs used, the memory items recalled, the models activated, the retrieval and ranking stages, the reasoning modes selected, the solver invoked and its result, the constraints applied, the experts consulted, the verifier or proof checker used, and the final expression path.

Deductive reasoning derives conclusions from rules. Inductive reasoning learns patterns from examples. Abductive reasoning ranks explanations. Analogical reasoning transfers relational structure. Causal reasoning examines interventions and dependency. Temporal and spatial reasoning preserve their own relation systems. Probabilistic reasoning uses fixed-point factor graphs where uncertainty is part of the problem. Decision reasoning evaluates actions and utility. Metacognitive reasoning monitors which reasoning mode is working. Counterfactual reasoning asks what would change under a different condition.

- Baseline
- Accept
- Exercise
- Review


QUESTION TO CARRY

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

ADOPTION

Adopt Loom 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 composed cognition path through adoption.

DECISION

Expand only what proved reusable

A catalogue can grow without every query becoming more expensive. Loom activates only the threads required by the weave. The organisation pays for the cognitive path that ran, not for every capability that was available.

A task that needs exact extraction and explanation. A task that mixes retrieval with formal constraints. A task that needs memory, causal reasoning, and verification. A task where a solver should replace a guess. A task that crosses domains without waking an entire monolith. See the weave Loom builds around it.

- 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 Loom 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 composed cognition path 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.

MECHANISM

Bring enough reality to make the session useful

Loom is built to stay light. Its helpers are small and efficient, so many jobs can run on ordinary computers, local servers, small devices, and controlled European systems rather than a distant data centre.

Loom has a catalogue of 528 experts, but they do not all speak at once. The relevant few can join when a question needs deeper knowledge or harder reasoning. The rest stay quiet.

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

QUESTION TO CARRY

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

THE INVITATION

Continue the Loom 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

DECISION

Turn the brochure into one bounded next step

Checking whether a set of rules can all be satisfied is not the same as writing a paragraph about them. The answer still arrives through one conversation. The work underneath uses the method that fits the question.

Loom is a sparse Mixture of Experts built on binary constraint satisfaction rather than a single weight tensor. A layered router selects a small active set per query. Inference is integer and bitwise, XNOR, AND, OR, and popcount. Outputs carry a full constraint trace, AION proof certificates, and are bit-identical under a fixed seed.

- 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?