

Product Overview

Eight commercial products / Current record 7 August 2026 / dweve.com/press

The Dweve Platform: Dweve B.V. is an Arnhem, Netherlands-based AI infrastructure company building explainable AI for Europe. The stack uses binary computation, constraint learning, and multi-agent orchestration to deliver AI that runs on standard hardware and can be traced, checked, and replayed where the declared execution contract permits it.

The commercial suite is eight products: Core (AI compute stack), Loom (constraint-learning runtime with 528 domain specialists), Nexus (multi-agent runtime), Spindle (AI knowledge governance), Aura (AI coding agent runtime), Mesh (Circular Compute infrastructure), Fabric (work-centred AI platform), and Charter (licensed company operating system). Fourteen open-source foundations are published alongside the suite, and Jacquard is documented separately as a research track.

AI COMPUTE STACK

Dweve Core

A complete AI compute stack for training and inference. Core keeps model authoring, training, inference, kernels, dispatch, and deployment in one coherent system rather than handing parts of it to other vendors.



Key Features

- More than 2,900 operations across 76 operation categories and 25 registered numerical format families
- Each implementation path is registered across nine dimensions: algorithm, datatype, width, packing, accumulator, backend, instruction set, dispatch policy, and determinism mode
- Scalar reference implementations plus target-specific paths across supported CPU, GPU, FPGA, browser, and deployment backends
- The runtime selects an implementation from the available machine and the policy attached to the work; paths are registered rather than inferred from a format label

Energy comparison: Dweve reports 96% lower energy for a stated binary workload against an FP16 GPU baseline. The published comparison includes its assumptions and is not independent certification.

CONSTRAINT-LEARNING RUNTIME

Dweve Loom

A constraint-learning runtime with a catalogue of 528 domain specialists. A task-shaped cognitive graph selects a path-specific set of specialists and keeps refusal and execution traces visible.



Domain specialists

- 528 domain specialists across 10 clusters: foundational (core reasoning, mathematics, science and research, code and systems, language) and specialised (multimodal, domain-specific industries, meta-cognitive, verification, transfer and adaptation)
- A 456-specialist production base sits alongside 72 expansion modules in training through the 57-layer evolutionary pipeline
- Clusters are independent. Adding specialists to one does not require retraining the others
- Approximately 3 billion domain-specific constraints per specialist. Conflicts surface as typed refusals, not soft warnings

Key Features

- Constraint learning first. Small bipolar language components render the result graph; they do not decide it
- PAP traces, integer and bitwise inference, and a CPU-first execution path
- Bit-identical replay under a fixed seed, compared to the last bit, in any environment
- Versions are content-addressed by weights, constraint set, and training corpus snapshot, so an upgrade is a deliberate, pinnable event

MULTI-AGENT RUNTIME

Dweve Nexus

A Python multi-agent runtime for work that needs more structure than a prompt loop. It separates agents, task lifecycle, workflows, memory, and capability-based routing into distinct components.



Key Features

- 6 subsystems, 31 perception extractors, and 8 reasoning modes
- Python builders, a DweveScript compiler, and a CLI as complementary interfaces to the same implementation
- Typed handovers using A2A-compatible types for messages, tasks, artefacts, capabilities, skills, roles, and statuses
- A typed event bus alongside point-to-point messages, plus a declared blackboard structure that agents read and write
- Model Context Protocol supported as one of the contract boundaries where a customer wants it
- The trace keeps each handover inspectable; authority over the outcome remains with the customer

Applications: Objectives that cross people, systems, rules, and approvals. Work is decomposed, authority assigned, failure contained, and every contribution returns to one accountable outcome.

AI KNOWLEDGE GOVERNANCE

Dweve Spindle

AI knowledge governance with provenance, evidence, conflict-aware retrieval, and reproducible datasets. Spindle turns sources into facts that can be defended, with an audit trail a regulator will read.



Key Features

- A graph of Canonical atoms with six-dimensional quality scoring and forward and backward lineage on every fact
- 47-type bias reports, with the bias engine running as part of certification rather than as an after-the-fact audit
- A reviewer can rebut a tag; the rebuttal lands in the chain, with no silent override
- Adding a category is a versioned change with an effective date. Old atoms keep their tags
- Scoped source classes for SharePoint, Confluence, data warehouses, ticketing, and collaboration platforms
- Subscribe to atom promotions, threshold breaches, bias flags, or conflict detections over webhooks, or query the state machine through the API

AI CODING AGENT RUNTIME

Dweve Aura

An AI coding agent runtime with a closed development loop: bounded actions, human steering, quality checks, replanning, and an operational record around the models your team chooses.



Key Features

- Per-tool permission levels, audit logging, PII filtering, and circuit-breaker protection. External tools are not trusted by default
- Agents, skills, plugins, and workflows are markdown and YAML on disk. Six built-in agents ship with Aura; project, user, and global tiers merge at startup
- Config resolves across four layers, built-in defaults, user level, project level, and environment variables, merged per key
- A single quality command runs formatting, linting, compilation, tests, doc-tests, and a security audit in sequence, reporting failures against diff context
- An MCP client over stdio JSON-RPC 2.0 registers external tools in the same registry, permission gate, and audit trail as the builtins
- Agents share a task board and a cross-agent work ledger that tracks file-level claims to prevent conflicts
- Ships as a single static binary with no runtime dependencies

Boundary: Aura acts inside declared permissions with a human steering the loop. It is not offered as a self-modifying or unattended system.

CIRCULAR COMPUTE

Dweve Mesh

Circular Compute infrastructure for sovereign distributed execution across controlled locations, with explicit workload and ownership boundaries. Idle capacity an organisation already owns is put to work instead of replaced.



Key Features

- Four deployment postures: Dweve's managed European public network, managed participation with customer-owned nodes, a private organisational or consortium Mesh, or air-gapped. Membership, governance, telemetry, settlement, and external connectivity stay inside the chosen perimeter
- libp2p autodiscovery with no central registry, Raft consensus for scheduling and cluster state, and Byzantine-robust aggregation that filters manipulated federated updates
- Federated learning as a workload class, with homomorphic encryption, secret sharing, and zero-knowledge proofs inside the scheduler, and an epsilon-delta differential-privacy budget tracked per round
- Post-quantum transport, using ML-KEM-768 for key exchange and ML-DSA-65 for signatures
- Three determinism modes per workload class: hard (bit-exact replay, integer-only hot paths, seeded randomness), seeded, and creative
- Jobs decompose into tasks across up to one hundred nodes by default, operator-tunable, over a unified gRPC and REST API
- Bronze to Platinum service-level tiers enforced in placement and settlement

Settlement, stated precisely: capacity has an owner and work has a cost, so contribution earns compute credits and consumption spends them, with escrow reserving credits before a job begins. Credits are non-transferable service-accounting units inside a trust domain. They are not listed on an exchange, not mined, not traded, and not promoted as investments, and there is no blockchain as an economic layer. The published compute-sharing terms state that the route is not a token scheme.

WORK-CENTRED AI PLATFORM

Dweve Fabric

A work-centred AI platform in which knowledge, models, agents, workflows, teams, and evidence remain attached to durable work rather than scattered across tabs.



Key Features

- Fabric begins like a familiar conversation, then keeps the useful parts: documents become reusable knowledge and good procedures become saved flows
- An answer can keep the sources and decisions that shaped it, so a shared result carries its evidence instead of a screenshot's conclusion
- Every surface, agents, workflows, knowledge, chat, models, budgets, and audit, is exposed as a first-class API over REST and WebSocket
- Workspace-scoped keys and versioned contracts, all writing one tenant audit chain
- No private files that were not chosen for inclusion. Customer data stays inside the EU

COMPANY OPERATING SYSTEM

Dweve Charter

A licensed company operating system for governed customer, policy, and operating workflows. Customers, money, people, delivery, support, knowledge, and AI actions sit on one operating graph.



Key Features

- One company record from first signal to renewal, so a person is not copied from app to app and retyped along the way
- Every domain engine reads and writes the same typed graph, keeping records, policy decisions, and event history linked
- One policy path covering both people and AI actions
- The built-in helper works from visible records, shows what it proposes in plain words, and waits for approval before anything changes
- Runs within an eligible licensed Dweve deployment, with no separate Charter charge. The licence covers operation; source code and modification rights are separate choices

PRESS

Contact and evidence

Press enquiries go to press@dweve.com, with a written first response within 24 hours. Put the deadline in the subject line for time-sensitive stories.



- Product architecture, methodology, limits, and release terms are published on the product pages at dweve.com/products
- Model records, evaluations, risk management, incidents, and data governance are published in the Trust Centre at dweve.com/trust
- Partnership and co-branding questions: partners@dweve.com. Legal and privacy questions: legal@dweve.com. Responsible disclosure: security@dweve.com