Dweve
Weaving Intelligence

ENGLISH EDITION · JULY 2026

MARKETING BROCHURE · 16 PAGES

Mesh: the product brochure

A standalone visual introduction to Mesh, 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

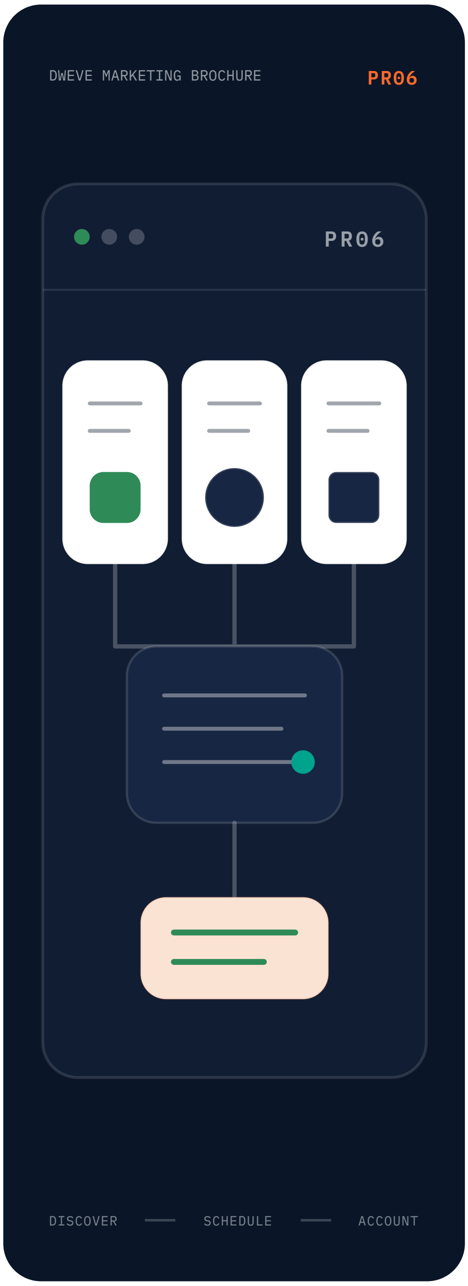

Decide whether Mesh belongs in the selected workflow.

PR06

THE PROPOSITION

Compute moves to the governed workload

Mesh orchestrates AI workloads across a peer-to-peer network of nodes, with automatic discovery, resource-aware scheduling, and privacy-preserving execution. Workloads include inference, training, retrieval, and agent workflows. Binary operations consume 96% less energy per operation than traditional float-based AI, making idle device participation sustainable. You can join as a managed participant consuming capacity, bring your own private nodes into the fabric, or run the mesh entirely inside your own perimeter.

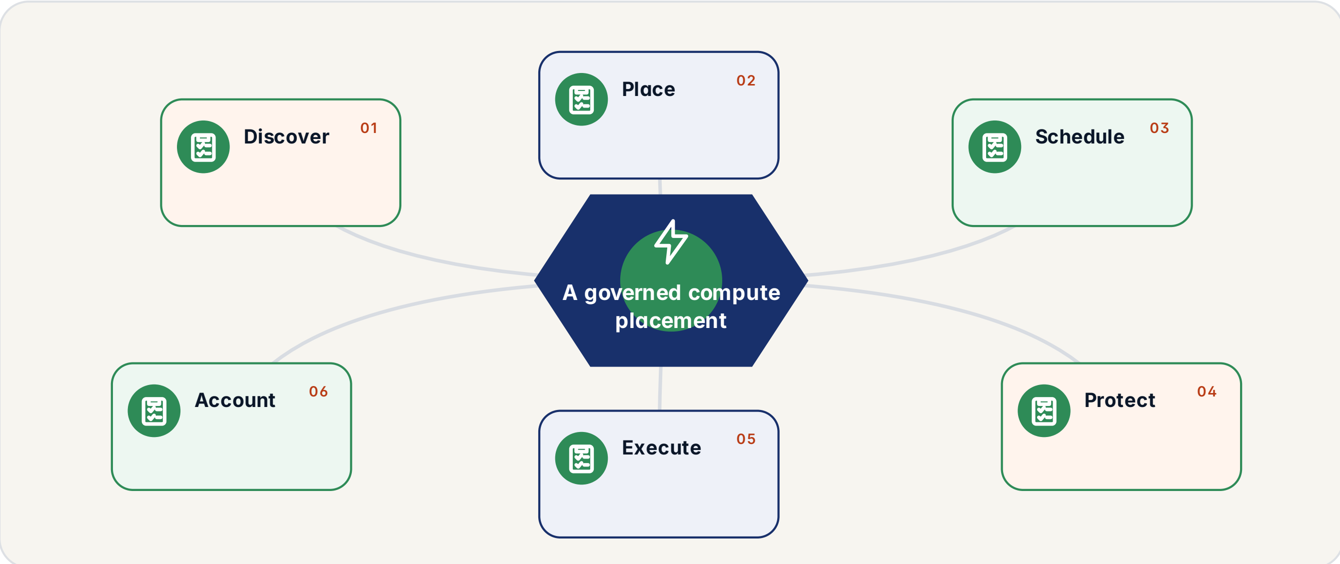

How to read this visual: follow a governed compute placement through the proposition.

MECHANISM

Read the promise against the operating reality

Mesh orchestrates AI workloads across a peer-to-peer network of nodes, with automatic discovery, resource-aware scheduling, and privacy-preserving execution. Workloads include inference, training, retrieval, and agent workflows. Binary operations consume 96% less energy per operation than traditional float-based AI, making idle device participation sustainable. You can join as a managed participant consuming capacity, bring your own private nodes into the fabric, or run the mesh entirely inside your own perimeter.

Mesh offers multiple scheduling algorithms, load-balancing strategies, and fault-tolerance modes so operators can tune for latency, throughput, or utilisation. Caching, connection pooling, and hardware-aware optimisation are handled below the workload. Binary operations consume 96 percent less energy per operation than traditional float-based AI, which is what makes idle-device participation sustainable. The defaults below are sane, and every one of them is tunable per task or per job.

- Discover
- Place
- Schedule
- Protect

QUESTION TO CARRY

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

USE AND FIT

Where Mesh can earn its place

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

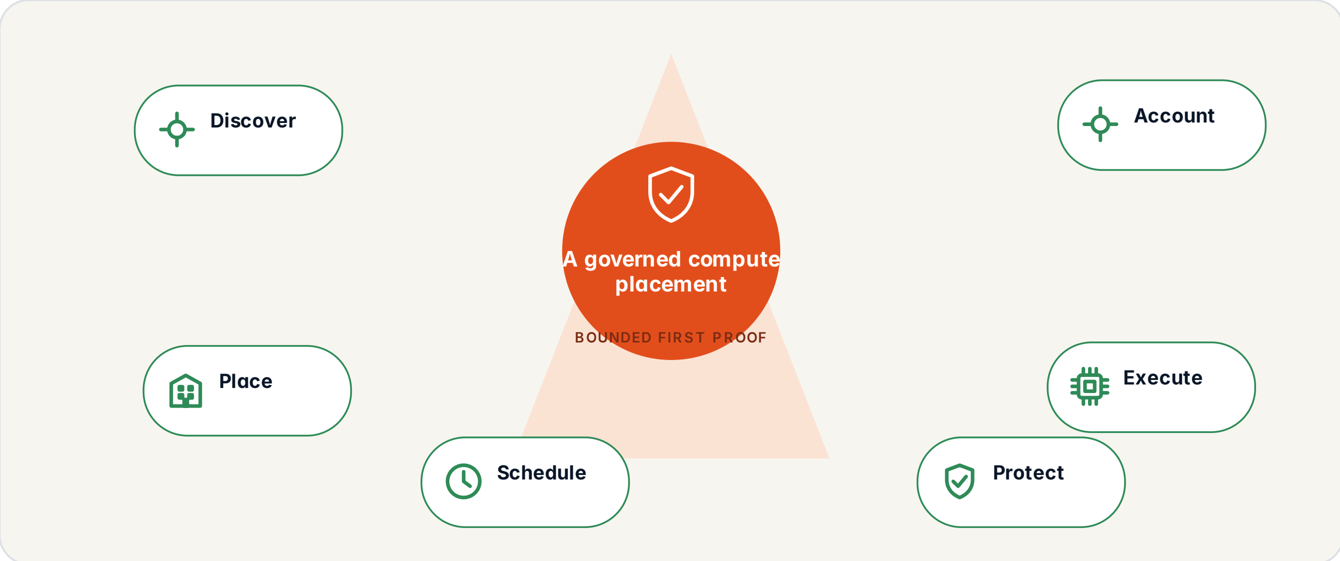

How to read this visual: follow a governed compute placement through use and fit.

DECISION

Separate useful scope from attractive distraction

Mesh discovers available CPU, GPU, NPU and FPGA capacity, matches it to suitable AI workloads, partitions work across heterogeneous machines and reallocates work when the available topology changes. It does not pretend computation is free: a processor doing useful work consumes more electricity than one doing nothing. The circular gain comes from using the hardware, baseline infrastructure, cooling, network, support and embodied resources already deployed. The machine already exists. Mesh gives its unused hours another purpose.

Mesh needs a settlement mechanism because capacity has an owner and work has a cost. Contribution earns compute credits; consumption spends them. Escrow reserves credits before a job begins and releases or refunds them according to the result. The credits are not listed on an exchange, not mined, not traded and not promoted as investments.

- Discover
- Place
- Schedule
- Protect

QUESTION TO CARRY

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

THE PATH

How Mesh moves from question to decision

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

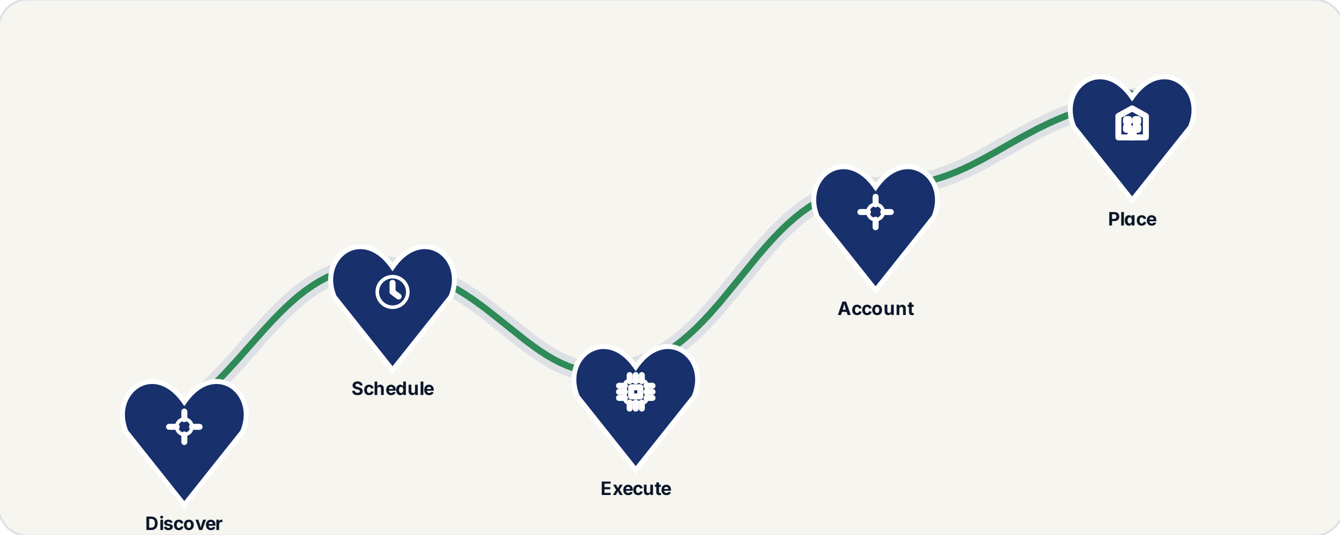

How to read this visual: follow a governed compute placement through the path.

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

This keeps topology mutation outside the request's logical execution contract. Brief disconnection linger and resume is the next optimisation layer, not a capability presented as shipped today: its purpose is to retain partition assignments and cache state for a configurable window before paying the cost of full recomputation. The claim that holds today: Mesh repartitions when capacity changes without silently changing the plan underneath an active request.

Dweve Mesh is a peer-to-peer execution fabric built in Rust. It discovers idle compute across heterogeneous hardware (CPU, GPU, NPU, FPGA), partitions AI workloads across it, and assembles results with cryptographic proof. Three inference backends: candle, llama.cpp, and dweve-native binary neural networks (96% less energy than float). No new data centers. No new power plants. Just the software to connect what is already powered on.

- Discover
- Schedule
- Execute
- Account

QUESTION TO CARRY

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

THE VALUE PICTURE

Mesh 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 compute placement through the value picture.

PRIMARY READER

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

The first conversation.

DECISION SUPPORTED

Decide whether Mesh belongs in the selected workflow.

The next useful step.

DECISION

Read the same outcome from several responsibilities

Mesh needs a settlement mechanism because capacity has an owner and work has a cost. Contribution earns compute credits; consumption spends them. Escrow reserves credits before a job begins and releases or refunds them according to the result. Rewards account for the work delivered, measured quality, demand and reliability. The credits are not listed on an exchange, not mined, not traded and not promoted as investments. Within a trust domain they move between accounts; exchange between separate domains uses explicitly governed agreements where that capability is enabled. The economic layer exists to make capacity allocation honest, not exciting.

Europe already contains millions of capable computers. Most of them operate alone: one machine waits while another organisation rents external capacity to perform work the first machine could have helped with. Mesh connects those spare hours into European AI infrastructure.

- Discover
- Schedule
- Execute
- Place

QUESTION TO CARRY

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

THE OPERATING MODEL

The control model around Mesh

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

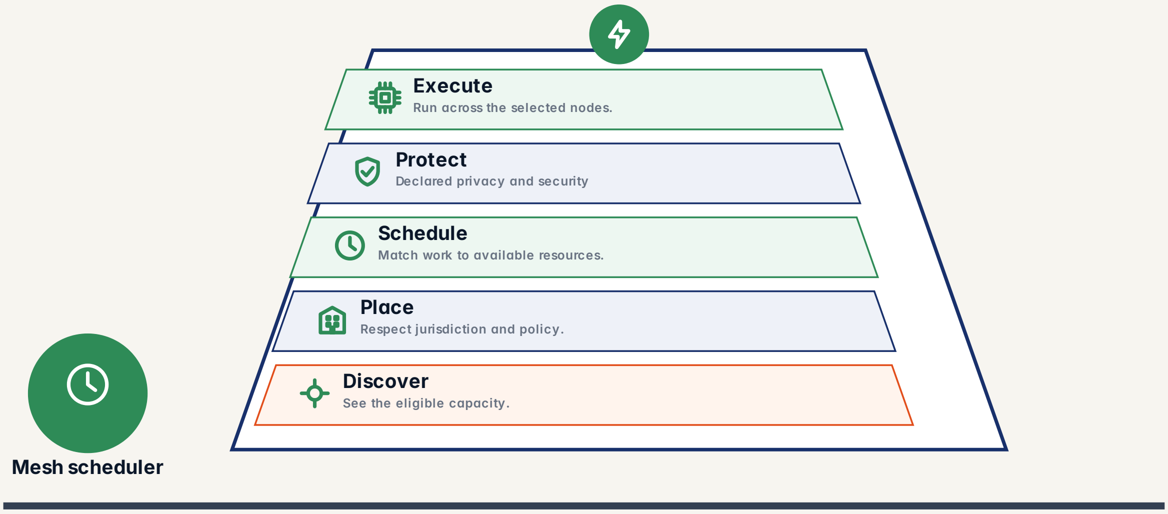

How to read this visual: follow a governed compute placement through the operating model.

MECHANISM

Name the controls behind the simple interface

Your organisation already owns processors, workstations, servers, edge devices and accelerators. They are powered, maintained, connected, secured and depreciating. Yet when a serious AI workload arrives, the default answer is still another cloud contract. Dweve Mesh changes the order of operations: it connects the capacity already available across your estate and turns it into a governed AI compute fabric. Use the estate first. Rent or build more only when the estate is genuinely full.

Dweve Mesh is a peer-to-peer AI execution fabric written in Rust. It discovers compute across heterogeneous hardware, partitions workloads across the nodes currently available, moves data and activations over encrypted connections, and assembles results with signed execution receipts. Nodes join. Nodes disappear. A laptop closes. A workstation owner returns. A GPU runs out of memory. A route dies halfway through a request. Mesh does not treat those events as surprising exceptions to an otherwise perfect cluster. They are the operating environment.

- Discover
- Place
- Schedule
- Protect

QUESTION TO CARRY

Who owns each control when the interface appears simple?

THE RESPONSIBILITY FLOW

People, Mesh, and evidence move together

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

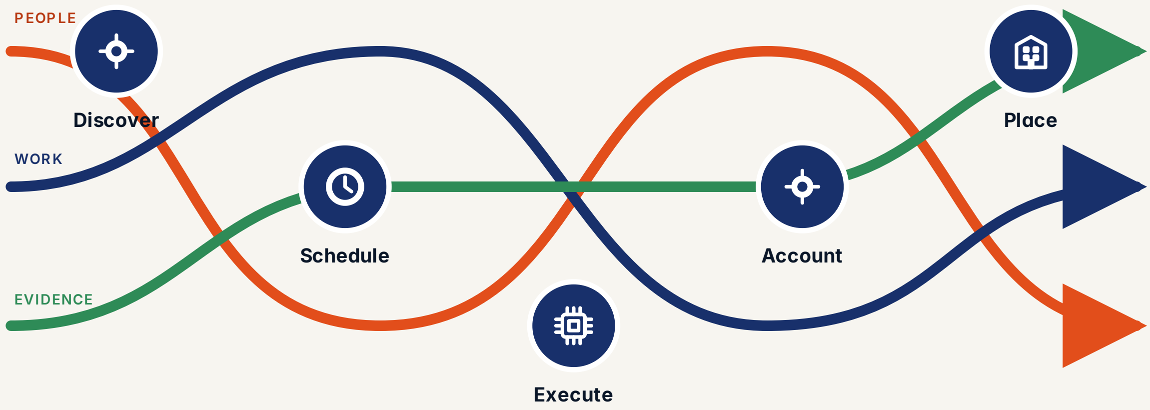

How to read this visual: follow a governed compute placement through the responsibility flow.

DECISION

Keep every hand-off attached to an owner

Mesh includes operational tooling that supports GDPR and ISO 27001 compliance work. Consent recorded by purpose. Data access and processing events logged. Portability requests producing machine readable exports. Deletion cascading through managed databases, storage, caches and sessions. Incident records preserving the sequence from detection through recovery. These controls do not grant automatic legal compliance or certification: they produce the evidence, enforcement points and repeatable procedures you need to demonstrate how your own policies were applied.

Mesh implements proposals for parameter changes, protocol upgrades, admission, expulsion and treasury allocation. Each proposal type defines its own threshold, quorum and voting period. Voting power combines reputation and stake under a cap. Delegation supports scope, expiration, chain resolution and cycle detection. Approved proposals execute through type specific handlers and produce audit entries. Governance is local to a trust domain: one private Mesh does not inherit another Mesh's constitution merely because the networks exchange work. Hierarchical representation and automated canary upgrade rollout remain design specifications and are labelled accordingly.

- Discover
- Schedule
- Execute
- Account

QUESTION TO CARRY

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

SOURCE TO RESULT

Mesh: from authoritative source to result

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

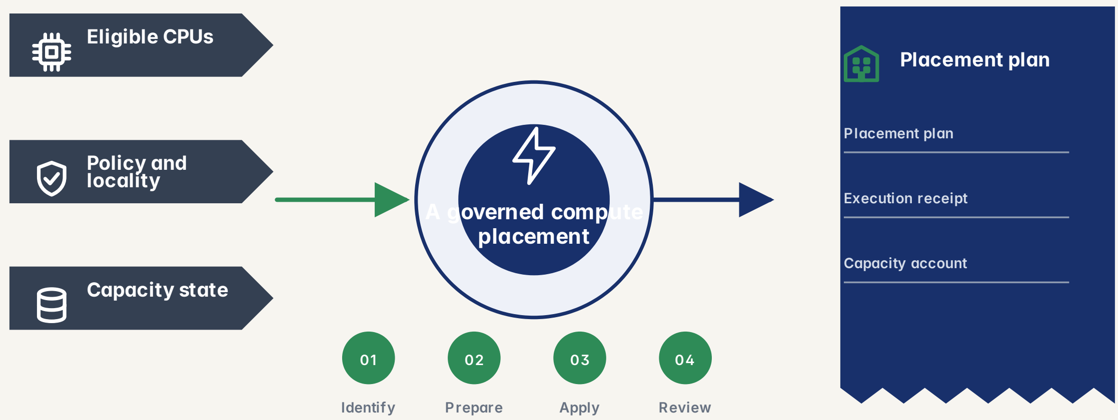

How to read this visual: follow a governed compute placement through source to result.

MECHANISM

Keep the basis visible as the work changes form

Some of the most valuable datasets cannot be centralised. Hospitals cannot casually pool patient records. Banks cannot merge customer transaction histories. Manufacturers cannot hand production telemetry to every consortium participant. Mesh makes federated learning a first class workload: the model is distributed to each participant, training happens inside that participant's own environment, raw records remain local, and only protected model updates return for aggregation. The protection mechanisms address different risks, and none is presented as magic. Together they provide a controlled way to learn across boundaries that raw data should not cross.

When AI runs on machines near you, responses arrive faster and your data travels less. Mesh routes work to European nodes by default, so your personal information never crosses borders you did not choose.

- Identify
- Prepare
- Apply
- Review

QUESTION TO CARRY

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

EVIDENCE PACKET

The evidence Mesh 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.

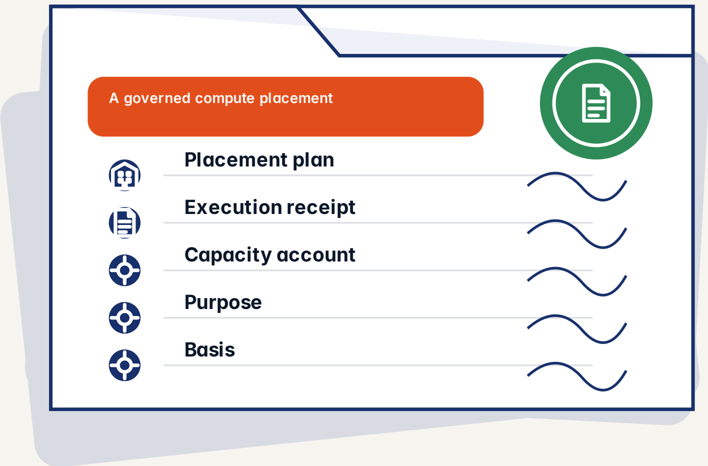

How to read this visual: follow a governed compute placement through evidence packet.

DECISION

Preserve the record needed by the next question

A trace on every decision means a real question comes back with the source it used, the rule it applied, the owner who signed, and a sealed record you can replay. Here is the shape of that evidence. The example is illustrative until your own data is connected.

Mesh exposes implemented REST and gRPC inference APIs, an OpenAI compatible chat endpoint, SSE streaming, operational endpoints for peers, storage, DHT, governance and WASM, and a basic CLI. The REST surface is currently broader than the CLI and gRPC surfaces; several extended commands and services remain planned. That asymmetry is documented plainly, because a distributed system should be able to explain both what exists and what has not earned the present tense yet.

- 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 Mesh

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

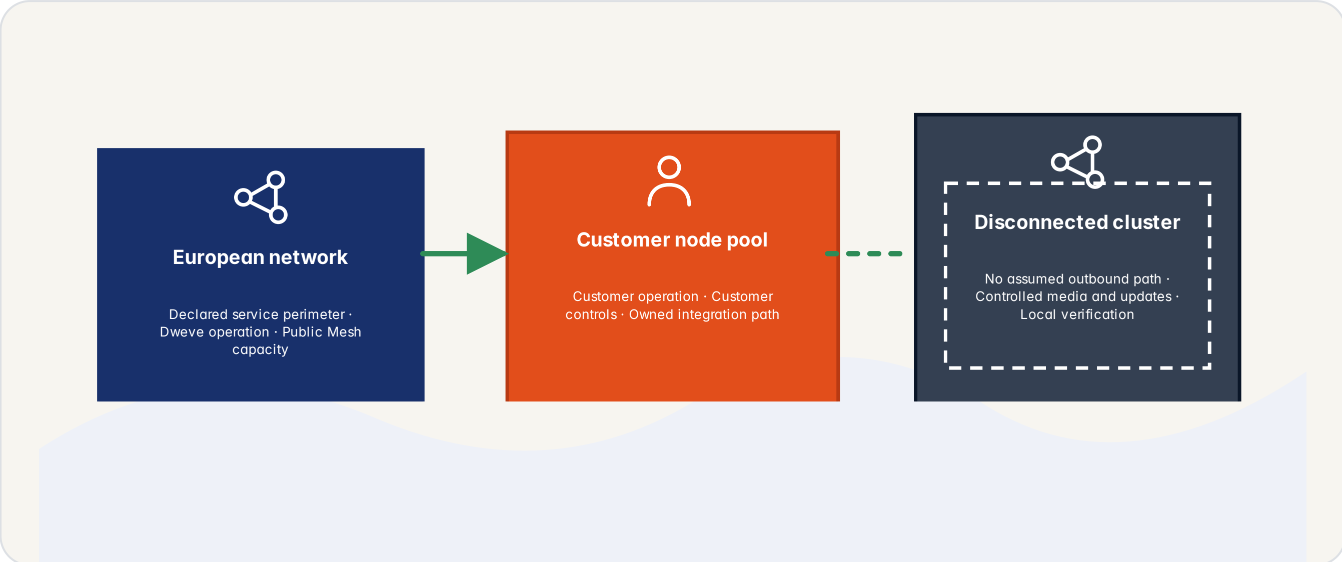

How to read this visual: follow a governed compute placement through deployment choice.

MECHANISM

Compare operating responsibility, not location alone

Managed customers participate through Dweve's public Mesh; Dweve operates that network and its capacity. A Mesh licence permits direct operation of a private or fully isolated Mesh on customer infrastructure. Scheduling, privacy, and governance boundaries remain explicit across both access paths.

Your enterprise runs thousands of machines. They boot at 9 AM, sit idle during meetings and lunch breaks, and stay on overnight for updates. The electricity is already paid for. The hardware is already depreciating. Dweve Mesh turns that idle capacity into productive compute, across CPU, GPU, NPU and FPGA, without new data centers or new power plants.

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

QUESTION TO CARRY

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

WHAT TO LOOK FOR

Signals that Mesh is earning trust

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

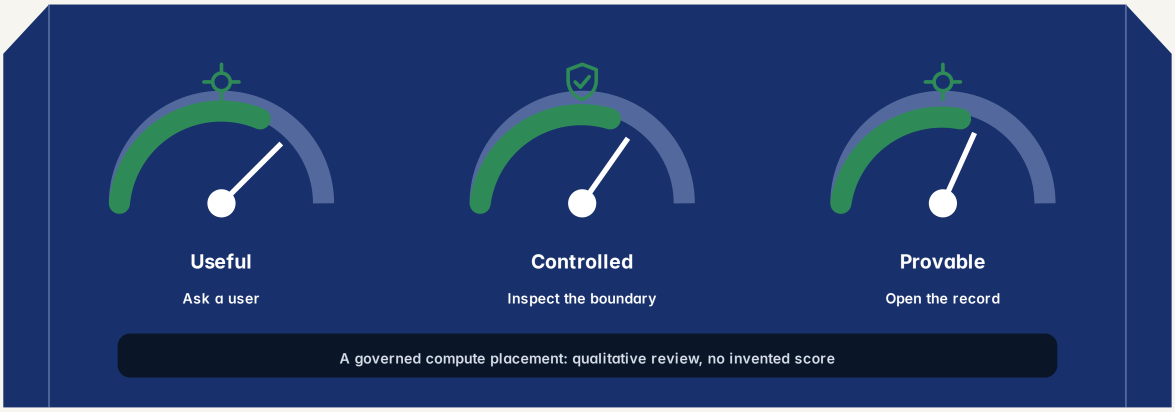

How to read this visual: follow a governed compute placement through what to look for.

DECISION

Use signals that can change the next decision

Local peers discover each other through mDNS. Wide area capability discovery uses a Kademlia DHT. Bootstrap nodes provide initial contact, not permanent control. Plumtree style gossip distributes service announcements, demand signals, governance results and peer state. Circuit relay and hole punching connect nodes behind NAT where direct routing is unavailable. Every peer is scored continuously on behaviour, reachability, latency, uptime and delivered capacity; low scorers are deprioritised, banned peers are excluded. The result is not a system with no infrastructure. It is a system in which normal execution, discovery and recovery do not depend on one central scheduler remaining alive.

Federated learning is not a feature toggle on Mesh; it is a supported workload class. Each participant trains on local data, shares only an encrypted update, and the mesh aggregates those updates into a new global model. The raw data never moves.

- Useful
- Controlled
- Provable
- Discover

QUESTION TO CARRY

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

HONEST LIMITS

Test how Mesh fails before the first success demo

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

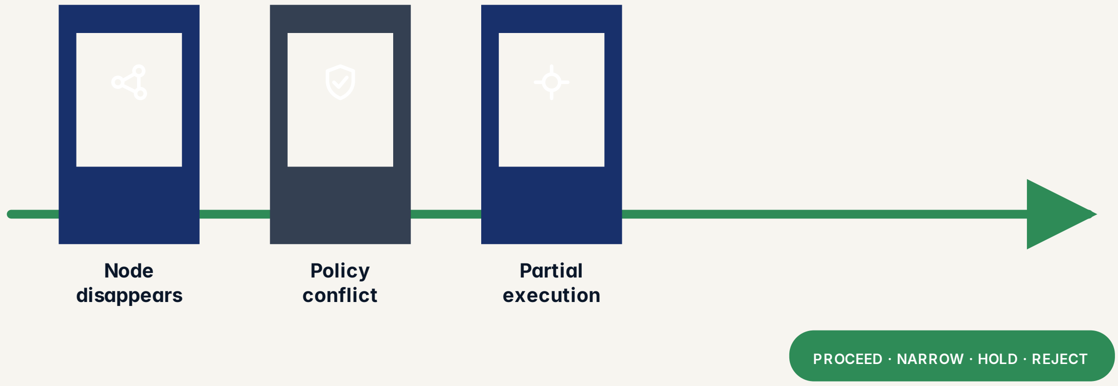

How to read this visual: follow a governed compute placement through honest limits.

MECHANISM

Design the failure before presenting the success

Circuit breakers exist at peer, service and DHT operation scope. Fault budgets connect error rates to operational state: a subsystem exceeding its budget moves into degraded operation rather than continuing to pretend it is healthy. The point is not that Mesh never fails. The point is that failure has a defined type, a bounded response and an inspectable history.

Mesh routes errors through one fault model. Every failure has a category, severity, retry classification and recovery strategy. Transient network failures use bounded exponential backoff; stale route failures retry immediately against another peer; validation, permission and cryptographic failures do not get retried until they accidentally pass.

- Node disappears
- Policy conflict
- Partial execution

QUESTION TO CARRY

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

ACCEPTANCE PATH

Close the Mesh proof with a decision

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

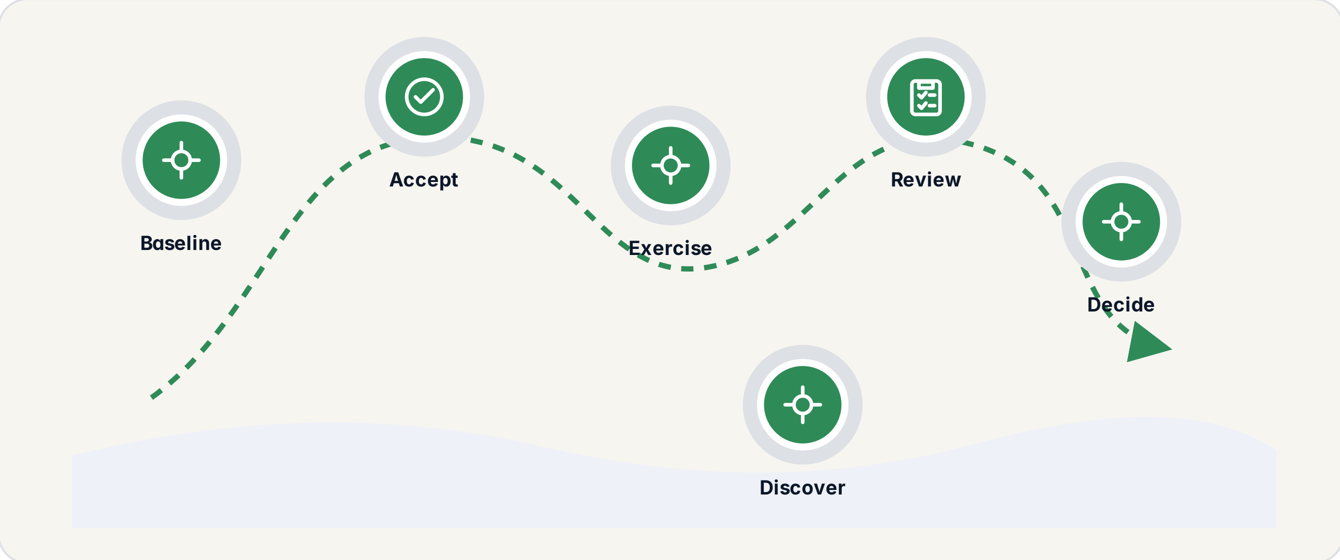

How to read this visual: follow a governed compute placement through acceptance path.

DECISION

Close the proof with an owned decision

Bring Mesh one workload and the machines already available to you. See which capacity can be recovered, how the workload is divided, where every part runs, what happens when a node disappears, and what evidence remains when the job is complete.

You do not need to understand distributed systems to use Mesh. Install the app, let it run when your machine is idle, and spend the credits you earn. Everything else happens automatically.

- Baseline
- Accept
- Exercise
- Review

QUESTION TO CARRY

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

ADOPTION

Adopt Mesh through owned choices

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

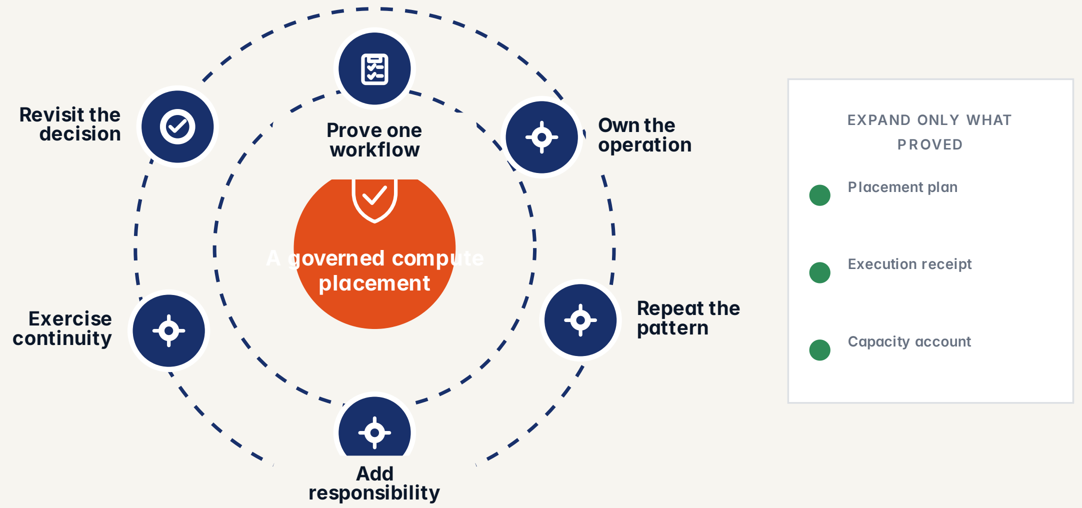

How to read this visual: follow a governed compute placement through adoption.

MECHANISM

Expand only what proved reusable

And each Mesh runs its own credit economy: Dweve issues and sells credits on the public network, while a private or consortium Mesh issues and governs its own, deciding internally how members purchase, allocate and settle them. Exchange between separate domains uses explicitly governed agreements where that capability is enabled. The economic layer exists to make capacity allocation honest, not exciting.

Compute is not a resource you extract, use, and discard. It is a cycle. The energy flowing through the chip is the raw material. Idle time is the waste. Mesh is the factory that converts that waste into value. Without new power plants. Without new data centers. Without waste.

- 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 Mesh working session

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

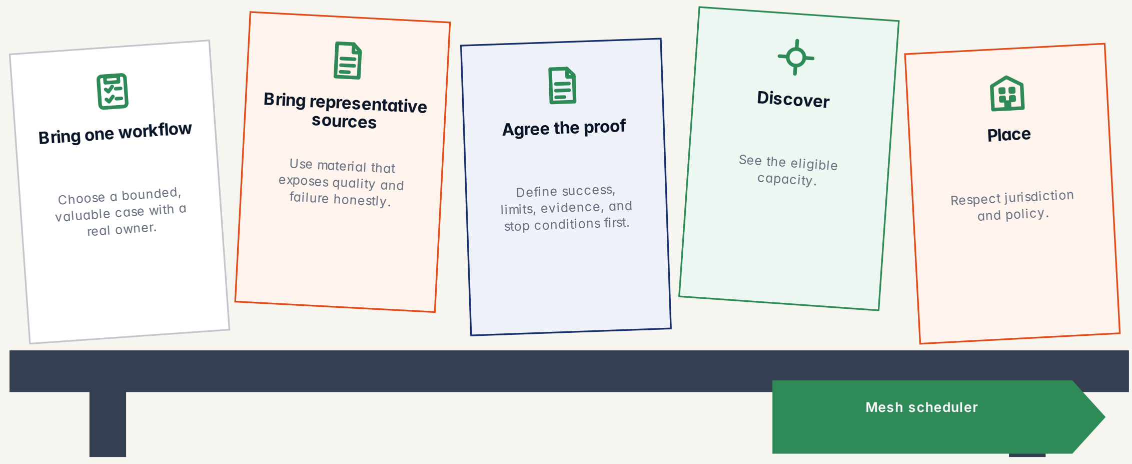

How to read this visual: follow a governed compute placement through what to prepare.

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

European energy costs make the economics compelling: if a company already pays for electricity to keep machines running, the marginal cost of running a Mesh job is near zero. Digital sovereignty is a political imperative: 78% of leaders consider it when selecting tech partners. Data centers are power-constrained: Europe cannot build its way out. Regulatory pressure is mounting: the EED came into full effect in 2025.

A hospital consortium runs their own private Mesh, keeps all patient data on German soil, and still accesses open-source models from the public network. No data leaves the jurisdiction. No US cloud provider is involved. Private meshes optionally peer to the public Dweve Mesh for overflow capacity and model access.

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

QUESTION TO CARRY

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

THE INVITATION

Continue the Mesh conversation

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

dweve.com/contact

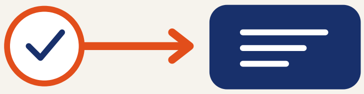

ONE MAINTAINED PUBLIC ROUTE

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

Mesh does not just use GPUs. It uses everything. A binary neural network runs on an NPU in a phone. A text embedding model runs on a CPU in an office desktop. A vision model runs on a GPU in a gaming PC. The hardware does not matter; Mesh abstracts it all away through the Selvedge runtime.

Every AI workload today has the same awkward shape: the data is yours, the regulator is yours, the users are yours, and the compute is someone else's. Mesh closes that gap with a European distributed network that brings compute to the data, not the other way around.

- 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?