EON UNIVERSAL

EON AI Ventures

Recognition → Twin

How a drawing + a few photos become a facility you can TRUST

The complete method — architecture, the novel core, what we proved, how we validate, and the forecast. Technical + commercial brief.

WHY THIS IS HARDER THAN IT LOOKS

A mesh is not a twin

✗ A generic 3D mesh

- One blurry blob
- No idea what a part IS
- No idea what connects
- No real size
- Cannot be operated

✓ An EON twin

- Every part identified + tagged
- Every pipe connected
- Real position, real scale
- A confidence on every value
- Operable: LOTO, training, maintenance

Buyers who say "we can 3D-model it ourselves" are picturing the mesh. The product is the twin — and the gap between them is the whole invention.

THE CORE INSIGHT

A facility is three separate truths

WHAT

Which parts exist & what each is (identity).

HOW

Which parts connect — the pipes (topology).

WHERE

Real position, size, depth in space (geometry).

No single input carries all three. Every failure we ever saw = asking one input for a truth it does not hold.

THE HONEST MATRIX

What each input can — and cannot — give

	WHAT	HOW (pipes)	WHERE (3D)
P&ID	✓	✓	✗ no coordinates
One photo	✓	~ partial	~ rough depth
Two photos (front+side)	✓	~ partial	✓ real depth
P&ID + photos	✓	✓	✓

The winning combination is P&ID (the brain: WHAT + HOW) + photos (the body: WHERE).

THE ROOT CAUSE

Why past renders came out as "spaghetti"

The seam carried positions with NO honesty flag. The builder could not tell a measured position from a schematic guess — so it improvised. Improvising on a flat, coordinate-free P&ID produces a fan of crossing pipes and floating parts.

The fix is NOT another layout tweak.

It is a confidence-governed contract: every value carries how-sure-we-are, and the builder is only allowed to improvise where the data admits it is weak.

WE STAND ON PROVEN SCIENCE

Prior art — each piece exists; the fusion is ours

1

P&ID digitisation

Symbol/line → graph. ~96–99% topology. (Digitize-PID, graph transformers.) → WHAT + HOW.

2

Scan-to-CAD

Scan2CAD / 9-DoF / FastCAD / SceneCAD: replace a rough scan with clean CAD at pose+scale. → your Tripo idea, published.

3

Scan-to-BIM pipes

PipeNet: don't trust scan pipes — rebuild from centerlines/graph. → route FRESH pipes.

4

Image-to-3D

TripoSR/Hunyuan: fast mesh, but scale-blind, unlabelled. → a scaffold, never the product.

Novelty = ingesting drawing + photos together, joining by tag, governing by confidence, anchoring scale to a library, self-validating. That combination is EON.

THE SYSTEM

Three-Signal Fusion — the pipeline

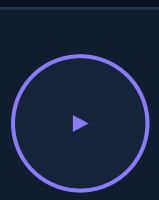

P&ID

identity + topology

Photos → Vision

labelled 3D boxes

Library

clean geometry + true scale

▼ **FUSE** (join by TAG · stamp every value with a confidence)

ONE model: per-part pose · size · orientation · 3 confidences (identity / topology / geometry)

▼ **BUILD** — place real models at measured pose · scale to library · route FRESH pipes · VALIDATE

Each stage takes only the truth its source is good at, and hands a clean, confidence-stamped contract to the next.

THE CLAIMED CORE

Five mechanisms that make it work

1 · Tag = universal join key

The drawing’s unique tags are handed to the vision model; every 3D box is labelled with a tag. Fusion becomes a deterministic join — no fragile geometric matching.

2 · Confidence-governed seam

Every value carries how-sure. The builder trusts geometry when high, falls back + labels "inferred" when low. Never a guess dressed as a fact.

3 · Library-anchored scale

Never trust the vision’s millimetres. Take relative pose from the photo; lock absolute size to the library’s true dimensions. One known part scales the whole rig.

4 · Fresh pipes from topology

Pipes are never lifted from scan or drawing — routed anew between placed parts from the authoritative connection graph. Owns no-crossings.

5 · Reverse-projection validation

Don’t build a noisy 3D to compare. Project the CLEAN twin back onto the photo and score overlap — ground truth, per part.

PROOF ON A REAL RIG

From a drawing + operator photos

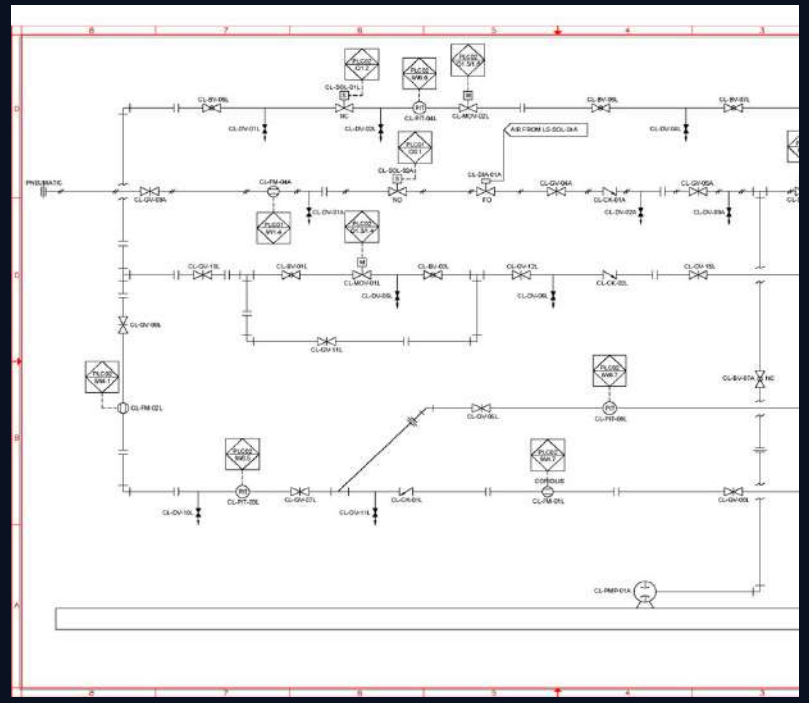

The P&ID (WHAT + HOW)

A photo of the real rig (WHERE)

137 parts + 133 pipes read from the drawing; 18 operator photos matched to their high-res originals by fingerprint; 21 of 22 left parts placed in real 3D.

THE RESULT

The reconstruction — measured 3D, not a flat guess

Real front-to-back depth across 8 distinct levels — the top-down view reads as a RECTANGLE, not a thin line. First pass came back flat; the method self-corrected using the side views.

PROVING IT, NOT ASSERTING IT

The validation stack — cheap tests first

1 · Scorecard	Counts parts, pipes, unconnected, confidence.	+ Objective, instant – No placement check
2 · Matched side-by-side	Photo vs model at the same angle.	+ Simple, intuitive – One angle at a time
3 · Reverse-projection %	Model onto the photo + overlap score.	+ A number; localises error – Needs the viewpoint
4 · Photo-3D vs model	Build a scan; compare 3D to 3D.	+ Most complete – Heavy; noisy reference

Validate the SKELETON before the builder commits, then re-check the finished twin. Run 1 always; 2 for the eye; 3 for a defensible number; 4 for deep audits.

WHAT TO EXPECT

Forecast of outcomes

Identity (WHAT)	95–98% parts correctly classed
Topology (HOW)	96–99% connections correct
Geometry — relative	right relative places; correct footprint
Geometry — metric	within a few % with a scale reference
Pipes	clean, connected, non-crossing

End-state: a P&ID + ~8 phone photos → a connected, labelled, operable, VALIDATED twin in minutes — vs. days-to-weeks of manual laser-scan-to-BIM at 10×+ the cost, and vs. a generic mesh that is unusable.

WHY THEY CAN’T JUST DO IT THEMSELVES

The moat — and the price it justifies

The intelligent library

Clean, dimensioned, operable models with engineering metadata — a years-long asset AND the scale anchor. Vision cannot manufacture it.

The fusion + confidence layer

Joining drawing + photo + scan by tag under a confidence-governed seam is the non-obvious core. Commodity parts don’t compose themselves.

Validation as a deliverable

A twin shipped WITH a per-part validation score is sellable to safety-critical ops. A bare mesh is not.

The honesty invariant

Refusing to present a guess as a measurement is exactly what makes it usable for LOTO, training, maintenance — the price-justifying use-cases.

The clearer the complexity is articulated and the more convincing the validated result, the weaker "we’ll build it ourselves" becomes — and the higher the defensible price.

THE PLAN

Roadmap

DONE	**Phase 0–1**	Multi-input intake · split confidence · synced side-by-side · scorecard · P&ID→pipes · 3D boxes from photos (proven on the real rig).
NEXT	**Phase 2**	Multi-view triangulation + library-anchored metric scale — kills the wrong-scale problem.
THEN	**Phase 3**	Optional image-to-3D scaffold + 9-DoF fusion for dense scenes.
THEN	**Phase 4**	Photos-only path · scale reference · reverse-projection + round-trip scoring.

EON AI Ventures

The one-sentence version

A P&ID says WHAT and HOW; photos say WHERE. We fuse them by tag, stamp every value with a confidence, place real components at real scale, route fresh pipes, and prove it by projecting the twin back onto the photo.

That is the invention — and the reason the result is worth paying for.