P&ID INTELLIGENCE PLATFORM — TECHNICAL BROCHURE

iCaptur.AI

P&ID intelligence for the engineered enterprise — reading piping & instrumentation diagrams as topology and delivering spec-driven, ERP-ready material take-offs.

AUGMENTING NATURAL INTELLIGENCE WITH ARTIFICIAL INTELLIGENCE

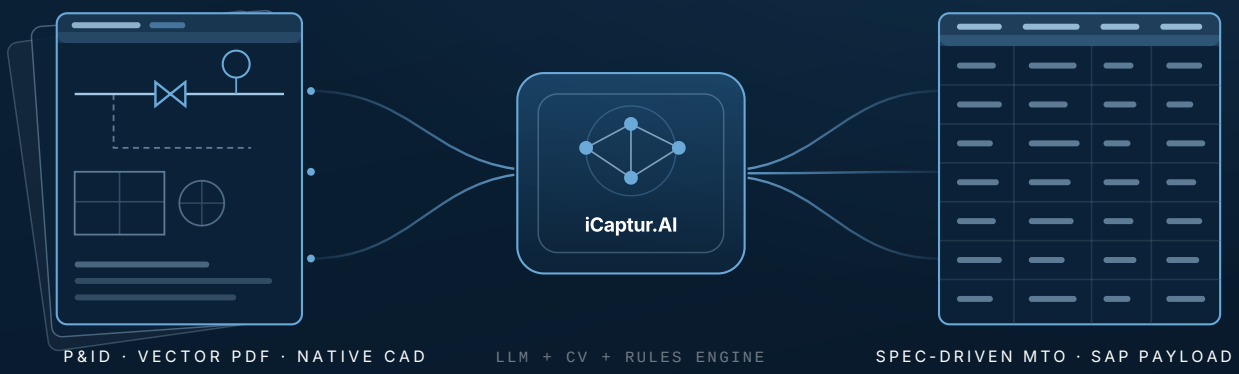

ISO/IEC 27001 · ISO/IEC 27701 · HIPAA · GDPR

Every P&ID already contains its material take-off. Reading it is the problem.

Industrial, refining and energy projects turn on thousands of Piping & Instrumentation Diagrams. Every one of them specifies, line by line, what has to be bought: the valves, reducers, blinds and instruments on each pipeline, and the flanges, gaskets and bolts each of those implies.

Today that specification is recovered by hand. Engineers identify symbols, trace line numbers across sheet boundaries and re-key quantities into a take-off — weeks of skilled time spent on transcription, with the errors surfacing late as procurement shortfalls, rush orders and budget deviations.

iCaptur.AI's P&ID Intelligence Platform reads the diagram the way an engineer does — as topology, not text. Vector geometry is resolved deterministically, symbols are classified against ISA-5.1 descriptors, and ASME B16.5 rules derive the mechanical joints, producing a spec-driven MTO with a traceable path from every output row back to the geometry that produced it.

The result is a take-off procurement can order against, in the SAP-compatible schema finance and supply chain already use — generated in minutes, reviewed by exception, and defensible to an auditor.

Where the hours go today

Manual take-off

Identify symbols | Trace line numbers | Re-key quantities

DAYS TO WEEKS · CROSS-HEADER ERRORS · NOT AUDITABLE

iCaptur.AI

Automated pipeline

Ingest | Resolve | Classify | Derive | Export | time saved

MINUTES · DETERMINISTIC · FULLY TRACEABLE

Same scope, different clock. On a 120-sheet P&ID package the platform compressed a three-week manual take-off into an eight-minute automated run (case study, page 05).

Where it lands in the business

Engineering	Take-off generated from the sheet instead of drafted from it; engineers review exceptions, not packages
Procurement	Bulk ordering against a complete, reconciled material list — joints included
Quality & audit	Every row carries its source sheet, coordinates and the specification rule applied
Systems	SAP-compatible MTO payloads and audit workbooks; nothing re-platformed

50M+ documents processed

500+ enterprise customers

99.9% uptime SLA

96% annual retention

72 net promoter score

MEASURED ON A 120-SHEET P&ID PACKAGE — CASE STUDY 01, PAGE 05

99.4%
Inline components auto-assigned to the correct pipeline line number

8 min
Material take-off across 120 P&ID sheets — down from three weeks manual

18k+
Mechanical joints derived without manual intervention

85%
Less downstream procurement reconciliation effort

Deterministic-hybrid extraction for piping and instrumentation diagrams

Industrial engineering, refining and energy infrastructure projects turn on thousands of Piping & Instrumentation Diagrams. Manual symbol identification, line routing and material take-off generation are among the most error-prone tasks in the entire project lifecycle — and the errors surface late, as schedule overruns and budget deviations.

Why P&IDs defeat generic document AI

A P&ID is not a page of text with a diagram on it. Meaning lives in topology: which line a valve sits on, which header a branch belongs to, where one line number stops and the next begins. Fragmented CAD exports break those paths into disconnected segments; intersecting utility headers invite *line-hijacking*, where a component is silently attributed to the wrong pipeline. OCR alone cannot resolve any of it.

The approach

The platform runs an enterprise-grade, deterministic-hybrid extraction pipeline. Non-raster PDF vector layer primitives are combined with topological ray-marching and rule-based ISA symbol classifiers, converting unstructured CAD exports into normalised, spec-driven engineering bills of materials and ERP-ready payloads — with a traceable path from every output row back to the geometry that produced it.

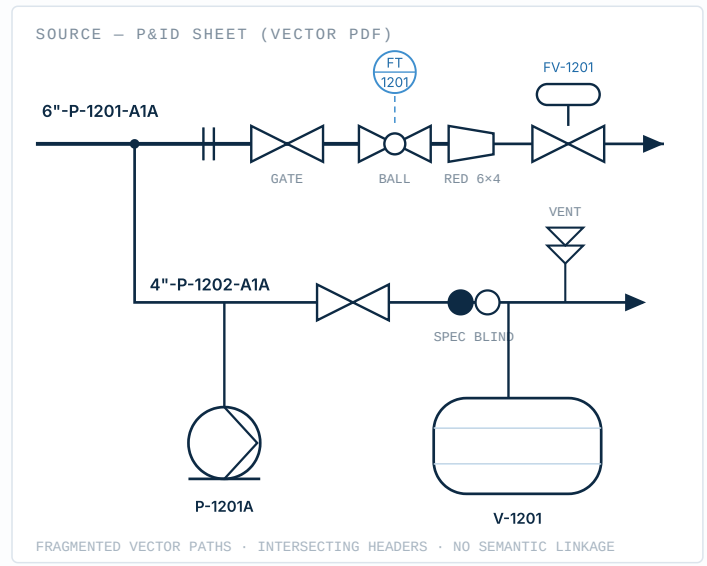

EXTRACT

OUTPUT — SPEC-DRIVEN MTO

LINE NUMBER	COMPONENT	SIZE	QTY
6"-P-1201-A1A	Gate valve	6"	1
6"-P-1201-A1A	Ball valve	6"	1
6"-P-1201-A1A	Control valve	4"	1
6"-P-1201-A1A	Reducer, conc.	6×4"	1
6"-P-1201-A1A	Flange, WN	6"	6
6"-P-1201-A1A	Gasket, spiral	6"	3
4"-P-1202-A1A	Gate valve	4"	1
4"-P-1202-A1A	Spectacle blind	4"	1
4"-P-1202-A1A	Vent tap-off	1"	1
4"-P-1202-A1A	Stud bolt set	M20	8

Derived joints added automatically **17 of 27 rows**

Figure 1 — from sheet to take-off. Flanges, gaskets and bolt sets shown in the output are *derived*, not drawn: the platform calculates the mechanical joints each flanged component implies and adds them to the take-off automatically.

Standards

ISA-5.1 symbol descriptors; ASME B16.5 flange and thickness tables

Input

Non-raster vector PDF and native CAD exports, batch-ingested per package

Method

Deterministic geometry resolution; no GPU-bound inference for baseline paths

Output

Audit-traceable workbooks and SAP-compatible MTO payloads

Four modules, each solving one failure mode

01 Non-raster vector primitive parsing layer

- Ingests native CAD and PDF layer models using high-performance vector extraction primitives, parsing discrete line segments, closed rectangles and multi-point cubic and linear curves.
- Bypasses heavy GPU-bound computer-vision and OCR inference for baseline geometry by mapping spatial coordinate vectors directly, at sub-point precision.

`VECTOR PRIMITIVES` `SUB-POINT PRECISION`

02 Axis-locked topological ray-marching engine

- Resolves pipeline multi-branch ambiguity and *line-hijacking* through an inline topological ray-marching engine.
- Combines axis-corridor spatial casting with deterministic gap-bridging to trace fragmented CAD vector paths back to their originating line-number tags, with no cross-header bleed.

`CORRIDOR CASTING` `GAP-BRIDGING`

03 ISA-5.1 descriptor-based symbol classification

- Uses an invariant symbol-matching architecture that normalises binary mask grids alongside geometric statistics — aspect ratios, convex hulls and moment descriptors.
- Classifies in-line components dynamically: gate, globe and ball valves; concentric and eccentric reducers; spectacle blinds; spacers; and vent and drain tap-offs.

`ISA-5.1` `INVARIANT MATCHING`

04 ASME-compliant spec-driven MTO & joint derivation

- Cross-references extracted line sizes and material classes against project thickness lookup matrices, assigning schedule ratings and wall thicknesses automatically.
- Derives mechanical joints in full — matching flanges, spiral-wound and ring-joint gaskets, and high-tensile stud-bolt quantities for every flanged component.

`ASME B16.5` `JOINT DERIVATION`

How module 02 defeats line-hijacking

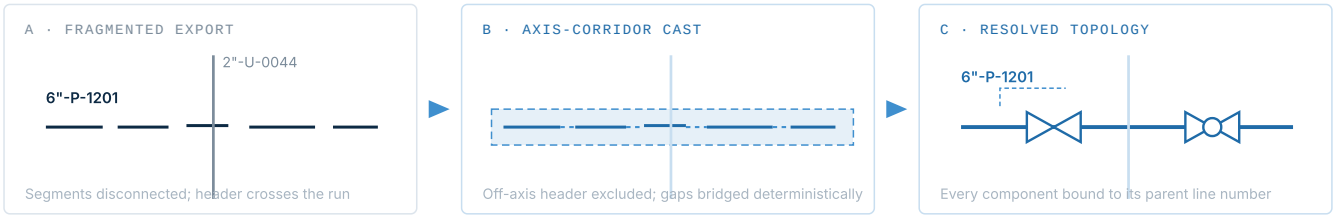

Measured on a 120-sheet package: **99.4%** of 4,500+ inline components assigned automatically to the correct parent line.

Why not simply throw a vision model at it

Pixel-level inference on a P&ID produces plausible answers with no provenance. When a take-off feeds procurement on a capital project, a confident wrong answer is more expensive than no answer at all — so baseline geometry is resolved from the vector layer, exactly, and models are reserved for symbol semantics.

What the auditor sees

Every output row carries its source sheet, its coordinates on that sheet, the classifier decision and the specification rule applied. Exceptions are surfaced rather than silently defaulted, so the review effort concentrates on the handful of genuinely ambiguous cases.

CASE STUDY 01 ENERGY & EPC · UPSTREAM HYDROCARBON PROCESSING

A 120-sheet multi-train facility take-off, executed in a single automated run

An international EPC contractor deployed the P&ID Intelligence Platform to produce a complete material take-off for a multi-train hydrocarbon processing facility. The engineering package ran to 120 complex P&ID sheets.

BACKGROUND

The package featured fragmented CAD vector exports, intersecting utility headers and dense instrumentation clusters — the exact conditions that break conventional extraction. Previous manual take-off attempts had suffered from inconsistent line-tag assignments, missed inline reducers, and wholesale omission of bulk joint components. Each omission surfaced later as a procurement shortfall.

IMPLEMENTATION

STEP 01

Batch ingestion & parsing

All 120 drawing sheets ingested concurrently, vector primitives parsed and title-block exclusion zones filtered — in under four minutes.

STEP 02

Topological line resolution

The axis-locked ray-marching engine mapped 4,500+ inline components — valves, reducers, blinds — to their parent pipeline line numbers.

STEP 03

Spec-driven rule application

ASME B16.5 thickness tables cross-referenced to generate 18,000+ derived mechanical joints — mating flanges, gaskets and bolt sets.

OUTCOME

Cross-header leakage was eliminated and unassigned component exceptions fell to near-zero. Standardised outputs exported directly into SAP-compatible MTO workbooks, so the procurement team worked from a reconciled list rather than rebuilding one.

MODULES DEPLOYED

VECTOR PRIMITIVE PARSING · RAY-MARCHING ENGINE · ISA-5.1 CLASSIFICATION · SPEC-DRIVEN MTO

What previous attempts got wrong

- Components attributed to the wrong parent line at header crossings
- Inline reducers missed on fragmented vector paths
- Bulk joints — flanges, gaskets, bolt sets — omitted entirely

What the engineering team reviewed

Rather than re-checking 120 sheets, reviewers worked a short exception queue: the components the classifier flagged as ambiguous, each with its source sheet and coordinates. Everything else carried a rule reference, so a sampled spot-check was enough to sign the package off.

MTO GENERATION TIME

8 minutes

→ down from **3 weeks** of manual engineering labour

99.4%
Automated line-number assignment across 4,500+ inline components

18k+
Mechanical joints derived without manual intervention

85%
Reduction in downstream procurement reconciliation effort

Output an auditor will accept, in the systems your teams already run

Extraction is only half the engagement. What determines whether automation survives is whether the result arrives in a shape procurement, QA and finance can use without rebuilding it — and whether every number in it can be traced back to the sheet it came from.

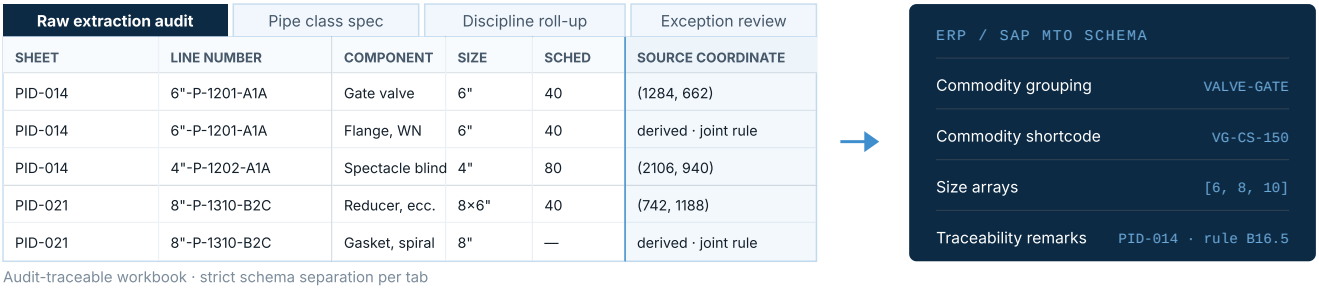

Figure 2 — deliverable architecture. Derived rows are labelled as derived and name the rule that created them, so a reviewer can distinguish what was read from what was calculated.

Structured multi-sheet workbooks

Fully formatted, audit-traceable spreadsheets with strict schema separation — raw extraction audits, pipe class specifications, disciplinary roll-ups and exception reviews each in their own sheet, never blended into a single flat export.

ERP & SAP data harmonisation

Output nodes map automatically into standardised enterprise MTO schema formats, including explicit commodity grouping, shortcodes, size arrays and traceability remarks — with PLM, EAM and digital-twin destinations supported alongside.

SECURITY & GOVERNANCE

Engineering drawings, patient records and legal case files are among the most sensitive documents an organisation holds. iCaptur.AI is built to be deployed against them — aligned to global regulatory frameworks, with confidentiality and secure handling designed into the pipeline rather than added at the perimeter.

- ISO/IEC 27001
- ISO/IEC 27701
- HIPAA
- GDPR
- 99.9% uptime SLA

Information security	ISO/IEC 27001-aligned management controls across processing and storage
Privacy management	ISO/IEC 27701 extension covering personal-data handling and retention
Health data	HIPAA-compliant record parsing with anonymisation and masking
Personal data	GDPR obligations supported, including data-subject handling
Auditability	Provenance recorded per output row: source sheet, coordinates, rule applied

05 — ENGAGEMENT

Start with your documents, not a demo

The fastest way to establish whether iCaptur.AI fits a workflow is to run it against a real package from that workflow. Engagements are structured to reach that evidence early.

STEP 01

Discovery

We review a representative sample of your drawings and documents, and agree what "correct" means for your standards.

STEP 02

Pilot

A bounded run on your own package, measured against your manual baseline — accuracy, exceptions and cycle time.

STEP 03

Integration

Rule profiles configured, outputs mapped to your ERP, PLM or QA systems, and review queues wired into existing approvals.

STEP 04

Scale

Rollout across disciplines and project sites, with throughput and exception rates monitored against the SLA.

Why iTech India

- Engineering depth as well as AI depth — the platform is built by people who understand what a pipe class, a datum reference and a first article inspection actually are.
- Deterministic-first architecture, so results are explainable and auditable rather than plausible.
- Enterprise security posture aligned to ISO/IEC 27001 and 27701, HIPAA and GDPR.
- Proven at scale: 50M+ documents processed, 500+ enterprise customers, 96% annual retention and a 72 NPS.

Bring us a hard package

The interesting test is not a clean, born-digital drawing set. It is the 1970s microfilm scan, the vendor package with three different title-block conventions, the P&ID where the utility header crosses six process lines. Those are the cases the platform was built around, and they are the ones worth putting in front of it.

P&ID take-off · Compliance screening · Raster-to-DXF · Ballooning & FAI · BOM consolidation

What a pilot needs from you

A representative package in whatever format it exists — native CAD, flattened PDF or scans — together with the standards it is held to and, where you have one, the manual result to measure against.

What comes back

The extracted dataset in your target schema, an exception report naming every case the platform would not assert on its own, and a like-for-like comparison against your baseline.

How it is scoped

Bounded to a fixed package and a defined acceptance measure, agreed before work starts — so the pilot answers a question rather than becoming a project.

iTech India Private Limited
icaptur.ai · itechindia.co

INDIA

2nd Floor, Unipunch Pride, #40,
2nd Main Road, Ambattur Industrial Estate,
Chennai 600 058, Tamil Nadu
+91 44 4385 8774 / 75
sales@itechindia.co

UNITED STATES

16803 Dallas Parkway, Suite 300,
Addison, TX 75001

+1 972 456 9479
sales@itechindia.co