



From Manual Take-Off to Traceable Intelligence

A Deterministic-Hybrid Approach to P&ID Extraction

EXECUTIVE SUMMARY

Engineering organizations have digitized their drawings for years. Native CAD, vector PDFs, and structured project data are standard infrastructure now. P&ID sets, specifications, and material lists exist as digital files instead of paper rolls in a drawer.

But here's the problem: while the format has changed, the extraction burden hasn't. A single project can involve thousands of P&IDs, and someone still must read each one by hand, tracing line numbers across sheet boundaries, identifying symbols, and re-keying quantities into a take-off. Going digital moved the drawings. It didn't remove the manual work of reading them.


That gap shows up downstream:

- Schedule overruns and budget deviations that trace back to a take-off error caught too late
- Procurement disputes over quantities that didn't match what was installed
- Rework triggered by a missed inline component or a mismatched line-tag assignment
- Weeks of skilled engineering time spent tracing and reconciling drawings instead of designing

This isn't a talent problem. It's the predictable result of asking a manual process to hold hundreds of relational judgments in mind across thousands of sheets, on a timeline measured in weeks.

The gap in P&ID extraction today is architectural, not effort based. Generic document AI reads text positionally and infers geometry from pixels, exactly the wrong model for a drawing where meaning lives in relationships: which line connects to which header, which component sits inline on which segment.

Real progress means resolving that topology first, deterministically, before any classification happens, and keeping every extracted value traceable back to the sheet it came from. That's where iCaptur's P&ID Intelligence Platform delivers value: turning a three-week manual take-off into a defensible, auditable output in minutes.





THE PROBLEM SPACE:

WHY P&ID EXTRACTION IS HARD

2.1 P&IDs as topology, not text


A P&ID is not a document to be read left to right. Its meaning lives in relationships, not positions, which line connects to which header, which component sits inline on which segment, where one line number stops and the next begins. A line number printed near a pipe run does not, by itself, tell you which run it belongs to. That has to be traced through the geometry. Any extraction approach that treats a P&ID as positional text will misassign a meaningful share of its output before it reaches classification.




2.2 Failure modes of manual take-off

Manual extraction fails in predictable ways:

- Inconsistent line-tag assignment on dense drawings or where multiple headers run close together
- Missed inline components, especially small symbols, unusual orientations, or reviewer fatigue on sheet 80 of a 120-sheet package
- Omitted bulk joints, since flanges, gaskets, and bolts often have to be inferred from spec rules rather than read directly off the page
- Cross-header errors at pipeline crossings, where a component gets attributed to the wrong parent line



2.3 Why generic OCR and vision-based document AI fails here




Document AI built for invoices or forms does not transfer to P&IDs:

- CAD exports are often fragmented at the vector level, with a single visual line composed of several disconnected segments that a generic parser has no reason to treat as one entity
- Intersecting utility headers create line-hijacking, where a tracing process follows the wrong line across a crossing point and silently reassigns a component to the wrong pipeline
- Pixel-level inference produces an answer with no record of how it was reached, so a wrong result looks identical to a right one until someone redoes the work manually
- OCR alone cannot resolve topology; it can read a label, but it cannot determine which pipe run that label belongs to

2.4 The cost asymmetry

A confidently wrong extraction is more expensive than no extraction at all. An output with no answer prompts a human review. An output with a wrong answer that looks plausible gets trusted, and the error moves downstream into procurement and fabrication, where it is far more costly to catch and correct than it would have been on the drawing itself.

Outcome: any extraction platform built for engineering use has to be designed around this asymmetry, not around raw throughput.



DESIGN PHILOSOPHY: DETERMINISTIC-HYBRID EXTRACTION

3.1 Determinism-first, not AI-first



This distinction is important because geometry provides the foundation for every downstream decision. If a component is attached to the wrong pipeline, accurate symbol classification cannot correct the resulting line assignment.

The platform treats baseline geometry as an engineering resolution problem rather than an inference problem. iCaptur uses deterministic parsing and rule-based logic to establish which vector segments connect, which path belongs to a pipeline, and where topology has been interrupted.

This distinction is important because geometry provides the foundation for every downstream decision. If a component is attached to the wrong pipeline, accurate symbol classification cannot correct the resulting line assignment.

The determinism-first architecture therefore provides:

- Repeatable geometry resolution from the same source input
- Explicit rules for connectivity and path resolution
- Consistent handling of fragmented vector geometry
- A stable foundation for subsequent classification and material derivation



3.2 Where machine learning is deliberately scoped

Machine learning is used where pattern recognition adds value: identifying what a resolved shape represents. It is not the primary mechanism for deciding which pipeline a component belongs to.

Once topology is resolved, the classifier identifies valves, reducers, blinds, spacers, and other components against defined symbol characteristics and engineering descriptors.

The distinction is:

- **Geometry:** vector structure and deterministic rules
- **Topology:** deterministic spatial logic
- **Symbols:** pattern recognition
- **Materials:** project and engineering rules
- **Provenance:** retained throughout

This keeps probabilistic inference within a bounded part of the workflow while allowing ambiguous results to be reviewed rather than silently propagated downstream.

3.3 Traceability as a first-class constraint

Traceability is treated as a property of the extraction pipeline rather than an audit report added after processing. Every output retains its relationship to the source drawing and the logic that produced it.

For extracted values, this includes the source sheet and coordinates. Derived values also retain the applied rule or specification. Every extracted value also carries a confidence score; values below the project's agreed threshold are routed to the exception review rather than counted in the total.

The provenance chain includes:

Source: drawing sheet and coordinates

Classification: component identified from resolved geometry

Rule: specification logic applied

Output: final MTO record and quantity

This allows reviewers to trace outputs without reconstructing the extraction process and focus human attention on exceptions.

PLATFORM ARCHITECTURE: THE FOUR CORE MODULES

4.1 Module 01 — Non-Raster Vector Primitive Parsing Layer



The parser produces a precise, structured representation of the drawing's native geometry, providing a reliable foundation for downstream topology and component resolution without relying on raster-based inference.

The first module reads the drawing's underlying digital structure rather than converting the P&ID into a raster image. This preserves the spatial information available in native CAD and vector PDF sources.

The parser converts native elements into a structured geometric representation:

- Discrete line segments representing pipe and connection paths
- Closed rectangles and other geometric primitives
- Linear and cubic multi-point curves
- Spatial coordinate vectors mapped at sub-point precision
- Title-block exclusion zones filtered before processing

This representation provides the geometric foundation for topology resolution. By operating directly on vector data, baseline geometry does not require GPU-bound vision or OCR inference, reducing unnecessary sources of uncertainty.

Where a sheet exists only as a scan, a separate vector-recovery step rebuilds the line work and isolates text before parsing, so the same parser and the same downstream rules apply. The confidence carried on each value reflects which route the sheet took.

Outcome: The parser produces a precise, structured representation of the drawing's native geometry, providing a reliable foundation for downstream topology and component resolution without relying on raster-based inference.



4.2 Module 02 — Axis-Locked Topological Ray-Marching Engine


Once vector primitives are parsed, the second module reconstructs drawing topology and determines which segments form continuous pipelines and which components belong to them.

It addresses line-hijacking, where tracing can follow an intersecting utility header and silently assign downstream components to the wrong line. A simple example: a 6-inch process line runs left to right and a 2-inch utility header crosses it vertically. At the crossing, a tracer that simply follows connected ink can turn onto the header and carry the process line's number onto every valve downstream of it. Nothing looks wrong on the page; the take-off has quietly moved those valves onto the wrong line.

The engine uses constrained spatial traversal:

- **Axis-corridor spatial casting** limits traversal to the expected pipeline direction
- **Deterministic gap-bridging** reconnects fragmented paths
- **Off-axis exclusion** prevents intersecting headers from becoming continuation paths
- **Topological resolution** establishes parent-line membership before classification

Outcome: components are assigned to their parent line without nearest-neighbour visual tracing. The reference deployment achieved 99.4% correct automated assignment across 4,500+ inline components.



4.3 Module 03 — ISA-5.1 Descriptor-Based Symbol Classification



With topology resolved, the third module identifies what each geometric object represents. Pattern recognition is applied here because symbol appearance can vary across drawing packages while following established engineering conventions. Before classification begins, the symbol library is built from the project's own legend sheet, so the reference shapes are the ones the drafting standard actually uses rather than a generic catalogue.

The classifier evaluates normalized geometric characteristics:

- Binary mask grid representations
- Aspect-ratio characteristics
- Convex hull information
- Geometric and moment descriptors
- Orientation and shape characteristics

The classification taxonomy includes:

- Gate, globe, and ball valves
- Concentric and eccentric reducers
- Spectacle blinds
- Spacers
- Vent and drain tap-offs

Classifications are mapped to ISA-5.1-aligned descriptor conventions, keeping the extracted component vocabulary familiar to engineering teams.

Outcome: symbol classification is separated from topology resolution, so the system can determine what a component is only after establishing where that component belongs in the drawing.

Where a symbol cannot be resolved to a single class — a flanged valve body whose type is not distinguishable from the symbol alone, for example — the component is listed for an engineer with its sheet and coordinates rather than assigned the most likely type.

4.4 Module 04 — ASME-Aligned Spec-Driven MTO & Joint Derivation



AP&ID is a schematic, not a scaled drawing. Roughly thirty per cent of a piping MTO — routed pipe lengths, fittings by run, supports — depends on the isometrics or the 3D model and cannot be derived from the P&ID at all.

The fourth module converts resolved and classified drawing information into a procurement-oriented material take-off. It derives materials implied by components, connections, line classes, and project specifications.

The rule engine cross-references extracted information against:

- Line size and material class
- Project thickness lookup matrices
- Schedule ratings and wall thickness
- Connection and flange requirements
- Gasket and stud-bolt requirements

For flanged components, it derives associated mechanical joints, including mating flanges, spiral-wound or ring-joint gaskets, and high-tensile stud bolts.

ASME B16.5-aligned rules are applied according to connection and line conditions rather than visible objects alone.

Outcome: the MTO represents both directly identified components and the materials required to complete their mechanical connections. In the reference deployment, more than 18,000 mechanical joints were derived without manual intervention.

What The P&ID Cannot Yield: a P&ID is a schematic, not a scaled drawing. Roughly thirty per cent of a piping MTO — routed pipe lengths, fittings by run, supports — depends on the isometrics or the 3D model and cannot be derived from the P&ID at all. The platform marks those items TBD with the reason rather than estimating them, so procurement sees the boundary of the number instead of a total that looks complete and is not.



FROM GEOMETRY TO DELIVERABLE: THE EXTRACTION PIPELINE END-TO-END

A P&ID enters the pipeline as raw vector geometry: line segments, curves, shapes, and text that do not yet form a structured engineering record. The four modules process this representation sequentially.

1. Resolution

Vector parsing establishes the drawing's geometric primitives. The axis-locked ray-marching engine reconstructs topology, connects fragmented paths, resolves gaps, and determines pipeline membership.

2. Recognition

Once topology is resolved, the classification layer identifies valves, reducers, blinds, spacers, and tap-offs using defined geometric descriptors and an ISA-5.1-aligned vocabulary. Classification therefore operates on resolved topology rather than page position.

3. Completion


Project and engineering specifications are applied to derive associated materials, including flanges, gaskets, and stud bolts. Derived materials are distinguished from information read directly from the drawing.

4. Handoff

The result is a structured, line-by-line MTO record retaining source information and identifying whether values were extracted or derived. The output can then move into workbook and SAP-compatible structures used by engineering, procurement, QA, and enterprise systems.

The four modules therefore operate as one dependency chain rather than as independent extraction tools. Geometry establishes topology. Topology establishes component context. Classification establishes component identity. Specification rules establish the associated procurement requirements. Traceability is maintained across each transition.





ENTERPRISE OUTPUT & INTEGRATION ARCHITECTURE

6.1 Structured multi-sheet workbook schema




Derived records are explicitly identified and retain the rule used to generate them. This allows reviewers to distinguish between information read directly from the drawing and information calculated from project specifications.

The extraction output is delivered as a structured, multi-sheet workbook that separates source data, specifications, roll-ups, and exceptions.

The workbook includes:

- **Raw extraction audit:** directly extracted values with source sheet and coordinate information
- **Pipe class specification:** line and specification information maintained separately from the component list
- **Discipline roll-ups:** aggregated material and component information for engineering review
- **Exception review:** items requiring human attention or additional validation

Derived records are explicitly identified and retain the rule used to generate them. This allows reviewers to distinguish between information read directly from the drawing and information calculated from project specifications.



6.2 ERP and SAP MTO data harmonization

The MTO is structured for downstream procurement and enterprise-system workflows rather than requiring another manual transformation step.

The output supports:

- Commodity grouping and standardized short codes
- Procurement-oriented size arrays
- Traceability remarks linking records to source drawings and applied rules
- SAP-compatible MTO workbook structures
- Downstream PLM, EAM, and digital-twin workflows

The objective is not to introduce another isolated data repository. It is to produce structured engineering information in formats that can move into the systems already used by procurement and project teams.

6.3 Audit trail design

Every output record retains its source sheet, spatial coordinates, classifier decision, and applied rule where the value was derived.

- **Source sheet:** drawing of origin
- **Coordinates:** spatial location
- **Classifier decision:** component identification
- **Rule applied:** derivation logic, where applicable

Exceptions are surfaced rather than silently defaulted, allowing review effort to focus on ambiguous cases.

- Flagged values can be traced to their source without re-running the take-off
- Standard records can be validated through sampled spot-checks



VALIDATION: MEASURED PERFORMANCE AT SCALE

Reference deployment: 120-sheet facility package

The platform was evaluated against a 120-sheet engineering package for a multi-train hydrocarbon processing facility containing fragmented CAD vector exports, intersecting utility headers, and dense instrumentation clusters. The package reflected conditions that can challenge conventional extraction, including incorrect line assignments, missed reducers, and omitted bulk joints.

The package was batch-ingested, topologically resolved, classified, and processed through the specification-driven MTO stage before export to SAP-compatible workbook structures.

Measured Results

Metric	Result
MTO generation time	8 minutes
Previous manual cycle	3 weeks
Inline components resolved	4,500+
Correct automated line assignment	99.4%
Mechanical joints derived	18,000+
Procurement reconciliation effort reduction	85%

Review-by-exception

Cross-header leakage was reduced to the exceptions tab: components that could not be assigned with sufficient basis were surfaced for review rather than defaulted, and the 99.4% figure above counts only assignments verified as correct against ground truth prepared by the client's engineers.

Ambiguous components were surfaced with source-sheet and coordinate information

Standard records retained rule references for sampled validation

This creates a review-by-exception model in which human engineering judgment remains part of the workflow, but is concentrated on cases where the automated pipeline does not have sufficient basis to assert an answer.

SECURITY, COMPLIANCE & GOVERNANCE

Engineering drawings are among the most sensitive documents an organization holds, and iCaptur is built to be deployed against them, with confidentiality and secure handling designed into the pipeline rather than added at the perimeter.

ISO/IEC 27001 alignment: information-security controls across relevant processing and storage activities

ISO/IEC 27701 alignment: privacy and personal-data management considerations where applicable

GDPR considerations: relevant data-handling and data-subject obligations where engineering document sets contain applicable personal data

HIPAA considerations: applicable only where a document workflow intersects with regulated health information, rather than as a requirement specific to P&ID processing

Provenance: source sheet, coordinates, classification information, and applied derivation rules retained with output records



DEPLOYMENT & ENGAGEMENT MODEL

The platform is intended to be evaluated against representative engineering packages rather than idealized demonstration data. The engagement model is structured as four bounded stages:

- **Discovery** — The team reviews representative drawings, source formats, project standards, specifications, and target output requirements. Acceptance criteria are established before processing begins.
- **Pilot** — A defined package is processed against an agreed baseline. The pilot measures relevant outcomes such as extraction accuracy, processing time, exception rate, and downstream reconciliation effort.
- **Integration** — Validated rule profiles and output schemas are connected to existing enterprise workflows, including ERP, SAP, PLM, QA, and approval processes where required.
- **Scale** — The workflow is expanded across drawing packages, disciplines, or project environments. Performance is monitored against agreed measures for throughput, exception handling, and output quality.

Pilot inputs

A pilot typically requires:

- Representative P&ID or engineering drawing packages
- Native CAD, vector PDF, or other source formats used by the client
- Applicable project specifications and standards
- Existing manual MTO or reference output, where available
- Defined acceptance criteria

The most useful pilot is not necessarily the cleanest package. Fragmented vector exports, dense instrumentation areas, multiple title-block conventions, and intersecting headers provide a

more meaningful test of the extraction architecture.

Pilot outputs

The engagement returns:

- Target-schema extracted dataset
- Exception report
- Source-level traceability
- Derived-material information
- Like-for-like comparison against the agreed baseline

Bounded-scope engagement philosophy

Each pilot is defined around a fixed document package, a defined scope, and an agreed acceptance measure. This keeps the engagement focused on an evidence-based question: whether the platform can deliver measurable improvement within the client's engineering workflow.

The pilot is therefore structured as a validation exercise before broader implementation and scale.



Each pilot is defined around a fixed document package, a defined scope, and an agreed acceptance measure. This keeps the engagement focused on an evidence-based question.

CONCLUSION:

REDEFINING THE P&ID-TO-PROCUREMENT PIPELINE

The deterministic-hybrid approach behind iCaptur reflects a specific view of where automation belongs in engineering workflows: geometry resolution is a question with a correct answer and should be handled deterministically, while symbol recognition is where pattern-based inference earns its place. Traceability isn't a feature added on top. It's a structural property of the pipeline itself, which is what makes the output something an engineering team can verify and an auditor can accept, rather than something that merely looks right.


Most of what a piping take-off needs is already on the P&ID; the part that is not — routed lengths, fittings by run, supports — belongs to the isometrics, and the platform says so rather than filling it in. What has made the rest expensive to recover is the manual effort of reading it, not any shortage of information on the sheet. Resolving topology first, classifying symbols second, and deriving what's implied third turns that recovery into minutes rather than weeks.

The same principles extend beyond P&IDs, to other document types and standards, without requiring a rebuild of the underlying pipeline.



ABOUT ITECH INDIA

iCaptur is built and delivered by the engineering practice of iTech India Private Limited — around fifty engineers, distinct from the company's software and AI services teams, including design-review and DFMEA specialists with more than twenty years each in industry. The platform is deployed as a layer that reads the documents a client already produces and returns files their systems already accept, hosted in the client's own environment where required. iTech India has operated for about two decades from Chennai, with centres in Tirunelveli, Villuppuram and Hyderabad and an office in Addison, Texas.





🌐 www.icaptur.ai

✉ info@icapstur.ai

☎ +91 44 4385 8774 / 75
+91 98405 95381