

Compliance Reasoning. Not Compliance Search.

This is the AI platform that reads **any** standards stack, reasons clause-by-clause, and cites every finding — built for engineering compliance and EPC reviews.

CHALLENGE

Engineering compliance is still done by hand — and the volume isn't going down. A single capital project carries thousands of datasheets, P&IDs, BODs, and specifications, each checked against owner standards, international codes (ASME, API, NACE, ISO), and project addenda by senior discipline engineers under deadline.

The first wave of AI tried to fix this with generic chatbots and similarity-based retrieval. They surface plausible-looking text. They don't **reason**. They miss negation, thresholds, paraphrase, and the addendum hierarchy that decides what a clause actually requires. They won't survive a serious technical review.

Engineering compliance needs a different category of AI.

THE COMPLIANCE BOTTLENECK

- 01 **Volume crushes review**
Thousands of clauses, hundreds of documents per package.
- 02 **The standards stack is layered**
Owner codes, international standards, and addenda that override.
- 03 **Misses are caught late**
A wrong parameter at design becomes scrapped fabrication.
- 04 **Senior engineers don't scale**
Most of their time is spent on lookup, not judgment.

THE SOLUTION

A new category: *Compliance Reasoning.*

Domain-engineered, not bolted-on.

The platform ingests your standards library — owner standards, international codes, project specs, and addenda — and applies it clause-by-clause against your engineering deliverables. The **reasoning** engine parses obligations, checks intent and thresholds, resolves cross-references, and returns severity-classified findings with full citation.

Every verdict links back to the source clause and the source quote. Confidence is reported separately from verdict — so reviewers know when to approve, when to escalate, and when to override. The result is an audit-ready compliance trail defensible to owners, auditors, and regulators.

WHY

01 / Standards-Aware

Reads **any** standards stack.

Owner standards, international codes (ASME, API, NACE, ISO), project specs, and addenda. Knows when an addendum overrides a standard.

02 / Reasoning, Not Similarity

Parses obligations, not paragraphs.

Clause-by-clause obligation parsing via the reasoning engine. Catches negation, thresholds, and paraphrase that similarity matching misses.

03 / Audit-Defensible

Cites the clause, every time.

Every finding links to source clause and source quote. Severity-classified. Confidence reported separately from verdict.

70–80%

reduction in review cycle time

10×

throughput per reviewer

60,000+

documents tested per engagement

Directional outcomes from active engagements — calibrate per opportunity.

— CAPABILITIES

The capability set. *Engineered for compliance.*

The platform is purpose-engineered for the way compliance reviews actually happen — from ingesting any standards stack to citing the clause behind every finding. Every capability below is part of the platform, not a roadmap promise.

- 01 Reads *any* Standards Stack**
Owner standards, international codes (ASME, API, NACE, ISO), project specs, and addenda — ingested as a project-aware corpus, not a generic search index.

02 Clause-Level Reasoning
The reasoning engine parses each clause as an obligation — checks intent, dimensions, exceptions, and cross-references. Reasoning, not keyword overlap or vector similarity.
- 03 Severity Classification *with Engineering Rationale***
Every finding tagged Critical, Major, Minor, or Observation — paired with the supporting clause, the engineering rationale, and the recommended corrective action.

04 Citation-First Output (*Confidence ≠ Verdict*)
Every verdict links to its source clause and source quote. Confidence is reported separately from verdict — so reviewers know when to approve, escalate, or override.
- 05 Addendum-Override Logic**
Project addenda override standards. The platform enforces the hierarchy automatically and shows reviewers exactly which clause version applied to a given finding.

06 Multi-Agent Orchestration
Built on LangGraph, Temporal, Mem0, and Qdrant for durable, long-running review jobs. ARIA, FORGE, and Planning Agent handle specialized review tasks under one orchestrator.
- 07 Reads Engineering Documents — *Not Just Text***
Datasheets, P&IDs, BODs, SPPs, drawings, and specifications — handled as engineering objects (tag numbers, dimensions, material specs), not as paragraphs to be searched.

08 Human-in-the-Loop with Learning Loop
Reviewers approve, override, or escalate every finding. Every decision is logged. Every override teaches the system. Compliance stays human-controlled, end-to-end.
- 09 Enterprise RBAC & Multi-Tenancy**
Platform Admin → Workspace Admin → Business User. Full tenant isolation with separate databases. SSO via Google, Microsoft, or your enterprise IdP.

10 Cloud-Native SaaS
The platform is delivered as a fully managed SaaS platform. No infrastructure to maintain, no version drift, no upgrade overhead. Always on the latest reasoning models and standards updates.
- 11 Open Integrations — *No Rip-and-Replace***
Plugs into your existing document estate — SharePoint, OneDrive, Aconex, Procore, Documentum, OpenText, SAP, Oracle ERP — via Model Context Protocol and REST APIs.

12 Audit-Ready Reports — *Cited & Defensible*
Cited Excel and PDF outputs that hold up under owner, auditor, and regulator scrutiny. Every line of every report traces to a clause, a quote, and a reviewer sign-off.

— A FINDING, ANATOMY OF

Every verdict ships with *the clause behind it.*

Severity, citation, source quote, confidence — separated for the reviewer's eye.

CRITICAL

DS-2401-COMP-V03 · Page 4

NON-COMPLIANT

Specified design pressure (78 barg) exceeds the maximum allowable working pressure for the selected flange rating.

CITED CLAUSE · ASME B16.5 § 2.3.1 (WITH PROJECT ADDENDUM P2-A14 OVERRIDE)

"The maximum allowable working pressure for Class 600 flanges in carbon steel at the rated temperature shall not exceed 75 barg..."

SEVERITY
Critical

VERDICT
Non-Compliant

CONFIDENCE
87%

Representative finding · clause references illustrative

— LET'S PROVE IT ON YOUR STANDARDS, YOUR DOCUMENTS

A scoped *2-week pilot.* Real outputs. No platform commitment.

We bring AI Compliance configured for your domain, an Engineering AI lead, and a solution architect — delivered as a managed SaaS environment.

You bring one representative standard, 8–10 sample documents, and one SME for two 90-minute working sessions.

You get a live prototype reviewing your real documents, findings benchmarked vs. your manual review, and a roadmap to production scope.

PILOT SCOPE

- 1 representative standard

- 8–10 sample documents

- 2 90-minute working sessions

- 1 SME from your side

- 1 working prototype, benchmarked

- 1 roadmap to production