

Tender Intelligence Pipeline

An AI system that reads **every public tender** across three national and EU procurement portals and tells us **which ones are worth bidding on** — each one filtered, scored on a six-dimension rubric, legally reviewed and explained. Engineered to run **entirely on local infrastructure**, so confidential tender documents never leave the company.

DOMAIN

Public procurement

DISCIPLINE

Agentic AI · Document intelligence

DEPLOYMENT

On-premise, privacy-first

CODEBASE

~8,000 lines · 5 stacks

EXECUTIVE SUMMARY

Winning public-sector IT work starts with a hard, unglamorous question: **of the hundreds of tenders published every month, which few are actually worth pursuing?** Answering it by hand is slow, inconsistent and easily crowded out by delivery work — so good opportunities are missed and effort is wasted qualifying ones that were never winnable.

Comtrade 360 built an end-to-end AI pipeline that automates that first-pass analyst work: it aggregates tenders from three portals, reads their documents, screens them for relevance, scores each survivor across six weighted business dimensions, runs a contract-risk review, and surfaces a ranked shortlist in a live dashboard. It runs against a **local language model by default** — no tender data is sent to a third-party cloud. And to de-risk the architecture, the **same solution was built five times**, once on each leading agent framework, so the technology choice rested on evidence rather than opinion.

3PROCUREMENT PORTALS
UNIFIED**6**WEIGHTED SCORING
DIMENSIONS**5**AGENT FRAMEWORKS BUILT &
COMPARED**100%**ON-PREM & PRIVATE BY
DEFAULT

01 The opportunity

Across the EU, a vast share of enterprise IT spend flows through **public procurement**. The notices are open and free to read — the difficulty is scale and friction. Opportunities are spread across separate national and EU portals, each with its own format, language and document conventions.

Qualifying a single tender means downloading a ZIP of PDFs and spreadsheets, reading the fine print, and judging fit against the company's technology stack, the people it can field, the deadline, the commercial terms and the contract risk. Multiply that by **hundreds of notices a month** and the honest first pass simply doesn't get done consistently.

The cost is twofold: **real opportunities slip through** unnoticed, while scarce pre-sales effort is spent qualifying tenders that were never a realistic win.

What "good" looks like

A system that reads **everything**, applies **one consistent standard** to every tender, shows its reasoning, and puts a ranked, decision-ready shortlist in front of the bid team — without any confidential document leaving the building.

Complete coverage

Consistent scoring

Explainable

Private

The three friction points it removes

- 1 Fragmentation** — three portals, one normalized feed.
- 2 Reading load** — documents parsed and summarized automatically.
- 3 Inconsistency** — every tender judged against the same rubric.

02 The solution — a six-stage pipeline

The pipeline performs the analyst's first pass end to end. Every expensive step is cached, so re-runs only reprocess what actually changed.

- 1 Fetch** — pulls every tender published since a given date from all three sources concurrently, normalizing them into one schema.
- 2 Analyse market trends** — classifies the **entire** fetched set (before any filtering) by solution area, buyer industry and required technology, feeding the Market Insights view.
- 3 Filter** — a two-stage relevance screen: a title triage drops the obviously-irrelevant, then a deep review of each survivor's full detail decides whether Comtrade 360 could realistically bid.
- 4 Summarize & review** — fetches each relevant tender's documents, scores six weighted dimensions with reason codes and red flags, runs a legal/contract review, and writes a Markdown executive brief with a Go / No-Go / Borderline call.
- 5 Shortlist** — keeps the strong matches (overall weighted score above threshold), ranked best-first.
- 6 Notify & export** — emails a digest of newly-added tenders and supports one-click SharePoint export of the gathered documents.

03 Inside the system

Roughly **8,000 lines** of application code, organized as a fleet of focused AI agents over a shared model client, three source connectors, a document-intelligence layer and a dependency-free web dashboard.

Connected data sources

PORTAL	REGION	SCOPE
eNarocanje	Slovenia	National public-procurement notices
eTenders	Ireland	National public-procurement notices
EU Funding & Tenders	European Union	SEDIA portal — EU-wide calls

The agent fleet

AGENT	RESPONSIBILITY	LOC
summarizer	6-dimension rubric scoring & brief	821
documents	Fetch / extract tender documents	710
lawyer	Contract & legal-risk review	457
trends	Market-trend classification	396
filter	Relevance triage + deep review	359
cvs	Talent-pool capability summary	296
chat	Grounded per-tender Q&A	191
exporter	SharePoint export gather step	61

Document Intelligence

Each shortlisted tender's attachments are fetched and text-extracted across every common format, then fed to the scorer as evidence.

- PDF
- DOCX
- DOC
- XLSX
- PPTX
- ZIP
- TXT / CSV

When a tender ships an oversized document set that exceeds the model's context, a **map-reduce evidence pass** summarizes the parts and reduces them into a single grounded brief — so large tenders score as reliably as small ones.

Durable caching

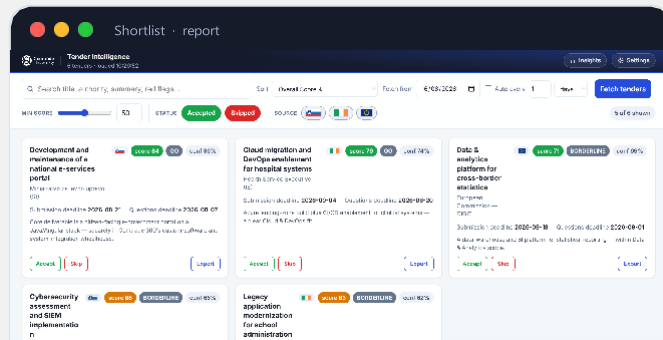
Every expensive per-tender result — filter verdicts, summaries, contract reviews, trend labels — is stored on disk, keyed by **tender ID + a hash of the model, prompt and schema**. Change a prompt and only the affected work re-runs; a repeat run over already-seen tenders is **nearly free**.

The dashboard

A **dependency-free** web app (no build step) turns the analysis into a working tool: a filterable, sortable **shortlist** with per-tender accept/skip decisions and a live "Fetch tenders" run button; a full **per-tender brief** with document upload, SharePoint export and a grounded chat panel; a **Market Insights** page of interactive demand charts; and a **Settings** page where the company profile, talent pool, scoring weights and model connection are all editable. An optional shared login gates the whole UI.

04 The product in action

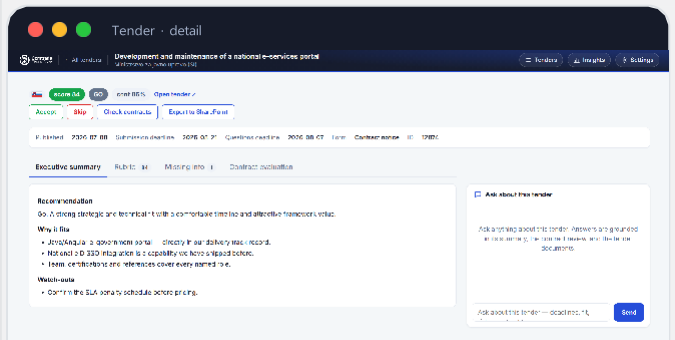
The same analysis, in the hands of the bid team. The dashboard is the delivered product — a live shortlist, per-tender briefs, market analytics and fully editable evaluation settings.



RANKED SHORTLIST

Every opportunity, scored and triaged

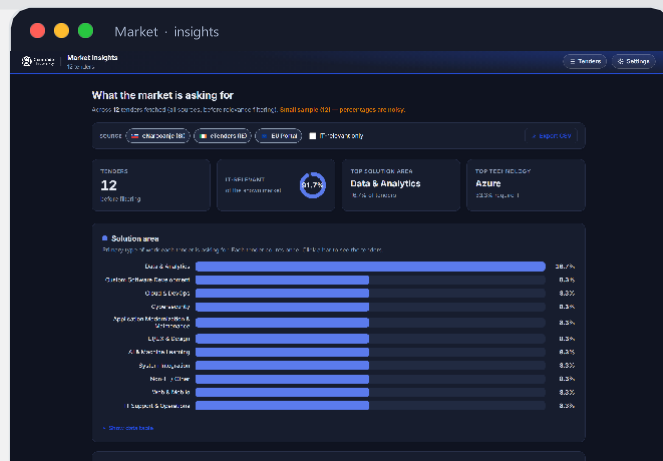
Score, Go / No-Go and confidence badges per tender, with source flags, filters and one-click accept / skip / export.



DECISION-READY BRIEF

Recommendation, rubric & grounded chat

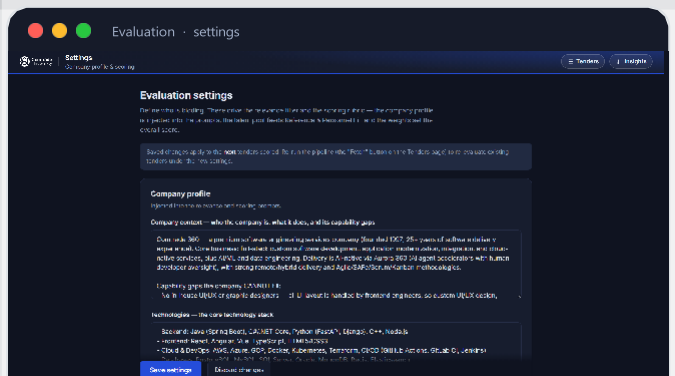
Executive summary with Go rationale and watch-outs, a full six-dimension rubric, contract review, and a chat grounded in the tender's own documents.



MARKET INSIGHTS

Where public-sector demand is heading

Solution-area, industry and technology demand across the full fetched set — filterable by source and IT-relevance, with CSV export.




EDITABLE, NOT HARD-CODED

Tune the model to your strategy

Company profile, talent pool, scoring weights and the LLM connection are all editable — the scoring reflects the real team, not fixed assumptions.

05 The scoring model

Every relevant tender is scored 0–100 on six business dimensions, each with a configurable weight. The overall score is the weighted average — and because the weights are **editable, not hard-coded**, the bid team tunes the model to its own strategy.



Default weights shown — adjustable from the Settings page; the overall score normalizes by their sum.

Grounded, not hand-wavy

Each dimension returns a numeric score, one or more **structured reason codes** (e.g. **TECH_STACK_MATCH**, **PERSONNEL_GAP**, **CERTIFICATE_GAP**), a written justification, and any **red flags** or missing-information notes. Reference & Personnel Fit is judged against an **anonymized capability summary of the real talent pool** — the model may not credit skills, seniority or certifications the pool does not actually list.

Calibrated to be conservative

The rubric is deliberately tuned to **default to the lower band**: a merely plausible fit is capped well below a clear one, and genuine uncertainty or missing documents pull the score down further. Hard guardrails encode real constraints — for example, a tender that demands in-house UI/UX design triggers a major red flag, because that capacity is known to be absent.

06 Market insights

Before any filtering, **every** fetched tender is classified — turning the raw feed into a live view of where public-sector demand is heading, filterable by source and IT-relevance, with click-through to the underlying tenders and CSV export.

12 solution areas

Custom Software · System Integration · Cloud & DevOps · Data & Analytics · AI & ML · Cybersecurity · Web & Mobile · IT Support · Hardware · UI/UX · Modernization · Non-IT.

12 buyer industries

Government · Healthcare · Education · Finance · Transport · Energy · Defense · Telecom · Manufacturing · Retail · Agriculture & Environment · Other.

30+ technologies

A steered vocabulary (Java, .NET, Python, React, AWS, Azure, Kubernetes, SAP, ...) keeps the same stack from being counted under three spellings; the top 15 are charted individually.

07 The engineering story — one solution, five frameworks

Rather than bet the product on a single vendor's agent framework, Comtrade 360 built the **same pipeline five times**. Every version shares identical agents, pipeline, dashboard and configuration — **only the thin model-client layer changes**. That turned an architectural guess into a controlled, apples-to-apples comparison, and proved the solution is neither framework- nor language-bound.

IMPLEMENTATION	FRAMEWORK	CLIENT	DEPS	WHAT IT DEMONSTRATED
Direct PY	Hand-rolled, no SDK	310	8	The baseline — total control, zero lock-in.
MAF PY	Microsoft Agent Framework	295	10	Fit with a Microsoft-centric enterprise stack.
LangGraph PY	LangGraph + LangChain	329	11	The most widely adopted open-source agent stack.
Pydantic AI PY	Pydantic AI	458	10	Type-safe, schema-first structured-output guarantees.
eve TS	Vercel eve	204	12	Full TypeScript rewrite — proof it isn't language-bound.

"Client" = lines in the single swappable model-client module. Everything else — ~8,000 lines of agents, pipeline and dashboard — is shared unchanged across all five.

Framework-independent

The valuable IP — agents, rubric, pipeline — is decoupled from any one SDK. Deploy on the stack a client already trusts, not the one a vendor prefers.

Evidence, not opinion

Framework choice was settled by building and comparing, not slideware. The same rigor is what a client's own architecture decisions get.

Model-agnostic too

Benchmarked across local and hosted models, standardizing on the one giving reliable structured output on commodity local hardware.

08 Two engineering details that make local models behave

Reasoning suppression

Modern "thinking" models can burn their entire token budget on hidden chain-of-thought and return an empty answer. The client forces **reasoning off** for the batch scoring calls, so the budget goes to the actual structured result.

Structured output

Every agent defines a strict JSON schema; the client wraps it in the model server's **response_format** envelope so results come back as **validated data**, not prose to be parsed — the backbone of consistent, comparable scoring.

09 Private by architecture

Tender documents are commercially sensitive, and the same pipeline is a template for far more sensitive internal data. So privacy is an architectural default, not an add-on:

- Runs against a **local language model** by default — **no data egress** to any third-party cloud.
- Pluggable to **any OpenAI-compatible API** (OpenAI, OpenRouter, vLLM, Ollama, ...) when a hosted model is preferred — the customer chooses where data lives.
- The model token is stored **server-side** and never sent back to the browser.
- An optional shared **login gate** covers every page and data API; ready to sit behind a **TLS-terminating proxy**.

The same system, the customer's choice

Whether a client mandates fully air-gapped inference or is happy with a hosted frontier model, the pipeline is identical — only the endpoint changes. That is the practical payoff of the framework- and model-agnostic design.

On-prem inference

No data egress

Optional cloud API

Login gate

TLS-ready

10 Outcomes

Hours → min

First-pass qualification, automated

Zero

Confidential docs sent to the cloud

Every tender

Screened against one rubric

5×

Portable across frameworks & languages

Complete coverage

Every published tender is read and classified before any human triage — nothing plausible is silently dropped.

Consistent & explainable

One rubric, reason codes and red flags on every tender — decisions the bid team can audit and defend.

Near-zero re-run cost

Durable caching means a daily run only pays for genuinely new tenders — the system is cheap to keep live.

The same problem, solved five ways — so we could choose the architecture on evidence, and deploy it wherever our clients' data needs to stay.

COMTRADE 360 · APPLIED AI ENGINEERING

Outcome figures illustrate the system's designed capabilities and privacy posture; qualitative claims reflect the delivered architecture. Let's talk about applying the same approach to your data — comtrade360.com.