



EXECUTIVE BRIEF

Agentic AI for Specialty P&C document intelligence

Turn ACORD forms, loss runs, SOVs, broker emails, and bordereaux into clean, validated, auditable data — with agents that classify, extract, validate, decide, and improve from every human correction.

Agentic pipelines

Purpose-built for P&C

Runs on any cloud

Any model

Free to download & use

Beyond OCR — an agentic pipeline

Most tools stop at OCR. InsightXtract understands the document, extracts to your schema, applies your rules, cross-checks systems of record, and routes to a human only when it should — pre-trained on the documents P&C actually receives.

01

Understand

Layout-aware extraction across ACORD forms, semi-structured loss runs and SOVs, and free-text broker emails — every field with a confidence score and a citation back to the page.

02

Decide

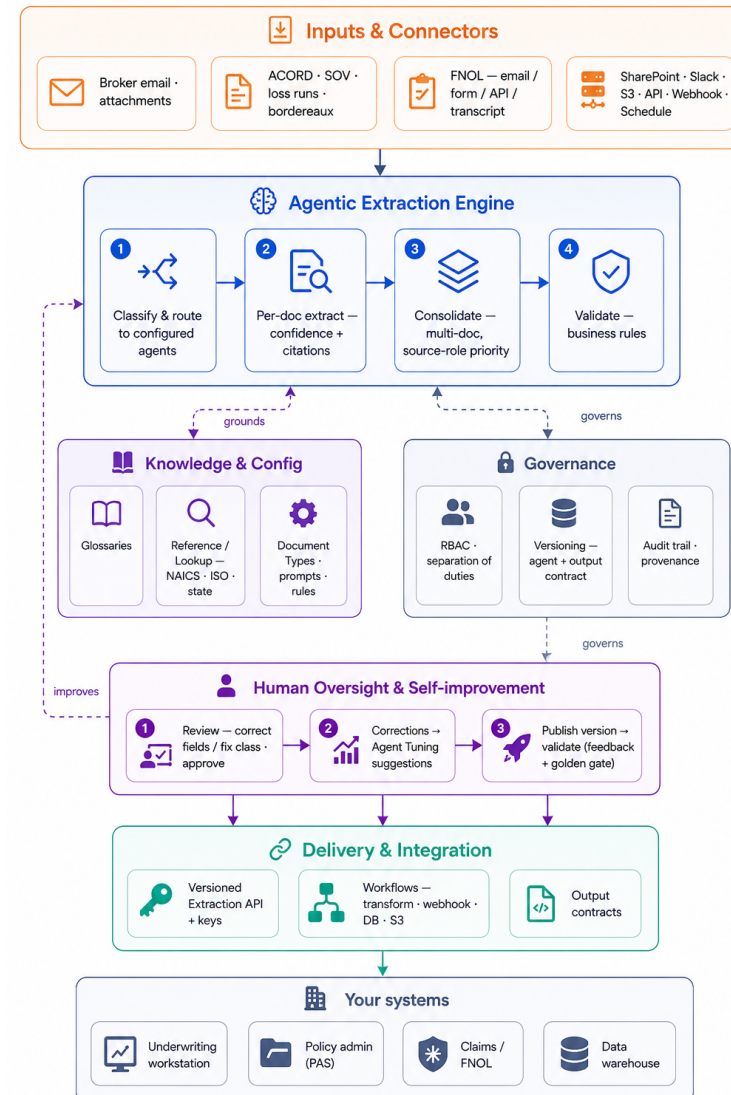
Configurable agents classify, validate against business rules, and consolidate multi-document submissions — straight-through where the data is clean, reviewer queue where it isn't.

03

Improve

Every human correction becomes a proposed prompt/rule change, validated against feedback and golden baselines, then published as a new agent version. Quality compounds.

High-level architecture



Agentic AI, end to end

Agentic multi-step extraction

Classify, plan, extract, validate, and reflect — with per-field confidence and reasoning, not single-shot OCR.

Multi-document consolidation

Each document routed to the right agent, then merged into one record by source-role priority (ACORD authoritative over email).

Human-in-the-loop review

Correct fields, fix classifications, approve or reject — every edit and decision captured with who and when.

Self-improving agents

Corrections become suggested prompt/rule changes, validated against feedback + golden baselines before promotion.

Provenance & versioning

Per-field bounding-box citations; agent and output-contract versions; a full, reproducible audit trail.

API, workflows & governance

Versioned Extraction API with scoped keys, visual delivery workflows, and role-based access with separation of duties.

High accuracy by design, not by luck

Accuracy comes from a robust agentic orchestration — reasoning at each step, a deterministic priority order when sources disagree, and an entity graph that ties the whole submission together — measured and gated before anything ships.

Reasoning, not single-shot

Agents classify, plan, extract, validate, and reflect. Each field carries a confidence score and a citation to the exact source location.

Priority-ordered resolution

When documents disagree, a declared source-of-truth order decides — ACORD over email, endorsement over base — deterministically, no guesswork.

Entity graph

Insured, policy, locations, vehicles, claimants, and coverages are linked across every document into one connected, de-duplicated record.

Domain knowledge built in

Glossaries and reference data — NAICS, ISO GL, SIC, state and currency codes — normalize values and flag invalid or unmapped codes.

Rules & validation

Business rules cross-check totals, dates, limits, and required fields; failures route to a reviewer instead of into your system of record.

Measured & gated

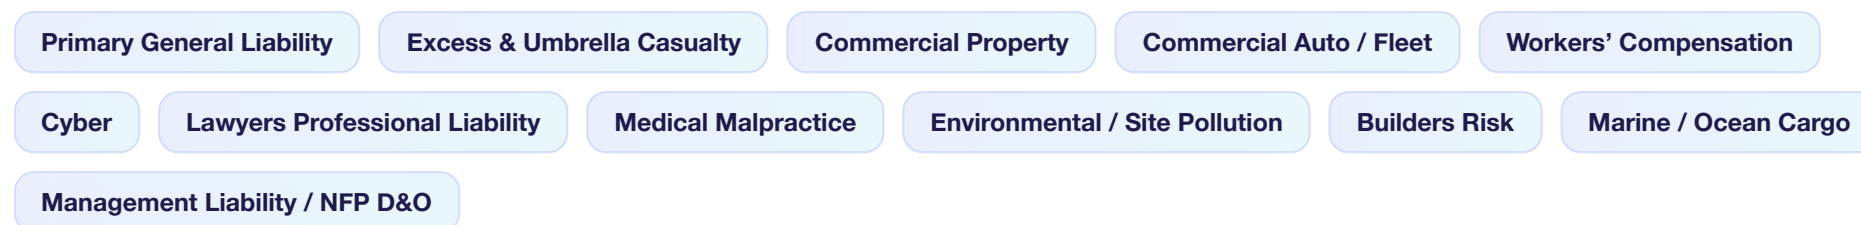
Golden datasets score field accuracy, STP rate, and F1 per document type; a new agent version must pass the gate before it's promoted.

COVERAGE

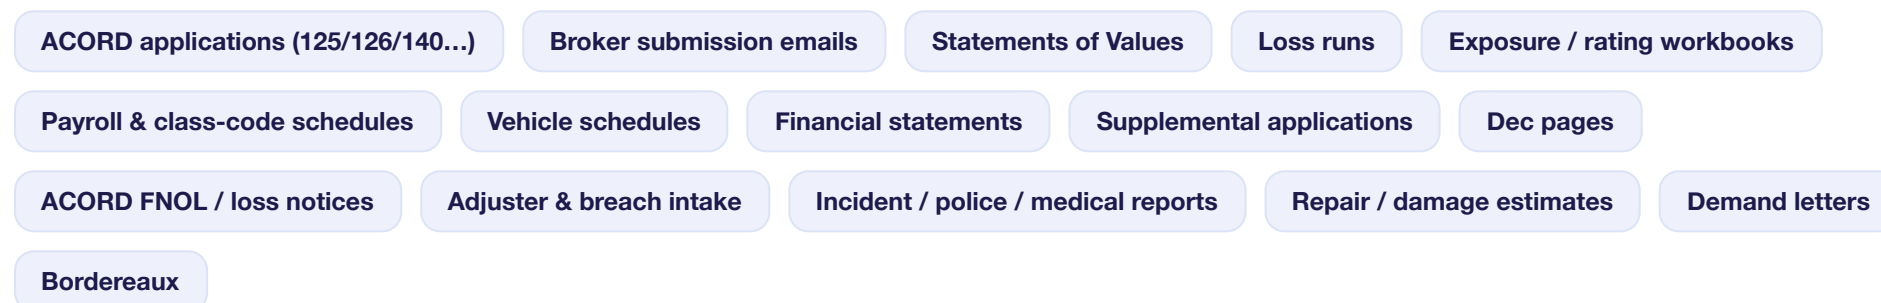
Covers the whole Specialty P&C book

Submission and claims intake across every major line — with the document types and broker templates that carriers, MGAs, brokers, and TPAs actually handle. New lines, document types, and templates are onboarded through configuration, not code.

LINES OF BUSINESS



DOCUMENT TYPES



BROKER TEMPLATES

Broker- and MGA-specific submission templates are mapped to your schema and onboarded without code — so a new producer's format is production-ready in days, not a project.

DEPLOY YOUR WAY

Runs anywhere. Any model. Yours to run.

Any cloud — or on-prem

Deploy on AWS, Azure, Google Cloud, or OCI, or fully on-premise. Everything runs inside your VPC and perimeter — your documents and data never leave your environment.

Any model

Anthropic Claude, OpenAI, and open-source or custom models — bring your own keys and route the best model per document type. No lock-in to a single provider.

Use the agents in your pipelines

Call any agent directly from your own pipelines through the versioned Extraction API and SDK — drop InsightXtract into what you already run, no rip-and-replace.

Two ways to run it

Try the hosted sandbox at insightxtract.com/console — just ask us for access — or download InsightXtract and install it in your own environment today.

How InsightXtract compares

Capability	Insight-Xtract	EXL Xtrakto.ai	Indico Data	Instabase	Hyper-science	Xtract.io	Xtracta	ABBYY	Forge AI	Aryn	Unstructured	Veryfi
Purpose-built for Specialty P&C	●	●	●	○	○	○	○	○	○	○	○	○
Agentic multi-step pipeline	●	◐	◐	●	◐	○	○	◐	◐	◐	○	○
Multi-document submission consolidation	●	◐	◐	◐	○	○	○	◐	○	◐	○	○
Domain knowledge (glossaries, NAICS/ISO)	●	●	◐	◐	○	◐	○	◐	○	○	○	○
Human-in-the-loop review + audit	●	●	●	●	●	◐	◐	●	◐	◐	○	◐
Corrections → agent self-tuning	●	◐	◐	◐	●	○	○	◐	○	○	○	○
Versioning + regression eval gate	●	◐	◐	◐	◐	○	○	◐	○	○	◐	○
Runs any model (Claude / OpenAI / custom)	●	◐	◐	◐	◐	○	○	◐	◐	●	●	○
Programmatic API + keys + usage	●	●	●	●	●	●	●	●	●	●	●	●
Role-based access + separation of duties	●	●	●	●	●	◐	◐	●	◐	◐	◐	◐
Deploy in your cloud / VPC (AWS·Azure·GCP·OCI)	●	◐	●	●	●	○	○	●	○	●	●	○

● Native / purpose-built ◐ Available / adaptable ○ Limited / general-purpose

Ratings reflect each vendor's general market positioning and public information, summarized for a P&C buying audience; several vendors are strong, mature platforms. This reflects focus for Specialty P&C, not a scorecard of overall quality. Product/company names are trademarks of their owners. Verify current details before external distribution.

The difference for Specialty P&C

P&C out of the box

Pre-trained on ACORD, loss runs, SOVs, and bordereaux — value in the first sprint, not after months of model training.

It gets better every day

Every correction becomes a validated, versioned improvement — quality compounds instead of plateauing.

Auditable and governed

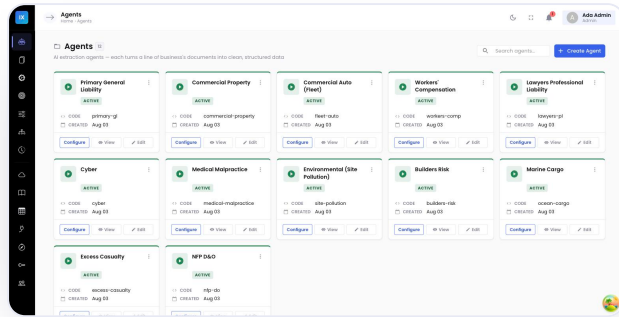
Provenance, versioning, RBAC, and a full trail — built for regulated carriers, MGAs, and TPAs.

Yours to run, no lock-in

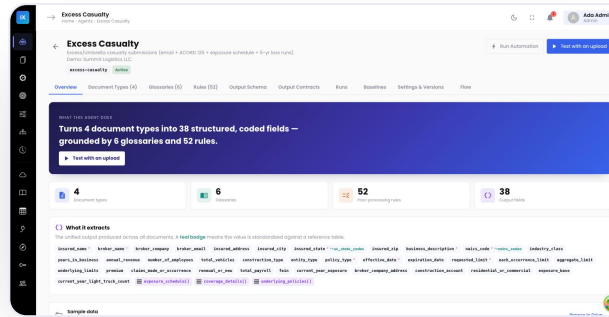
Any cloud, any model, callable by API — deployed inside your perimeter and dropped into the pipelines you already run.

IN THE CONSOLE

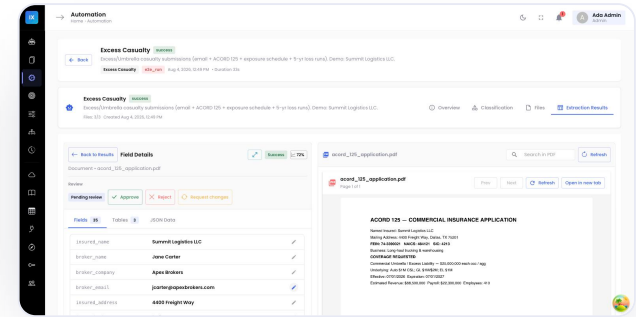
See it in the product



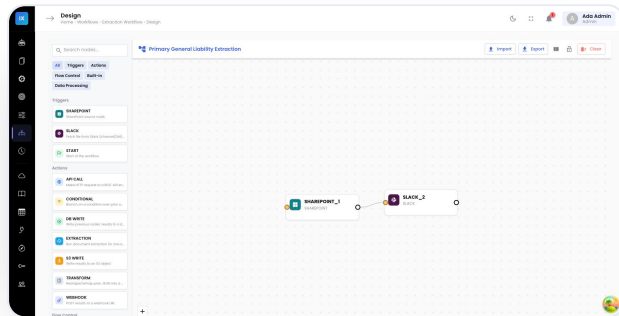
Agents. One extraction agent per P&C line of business.



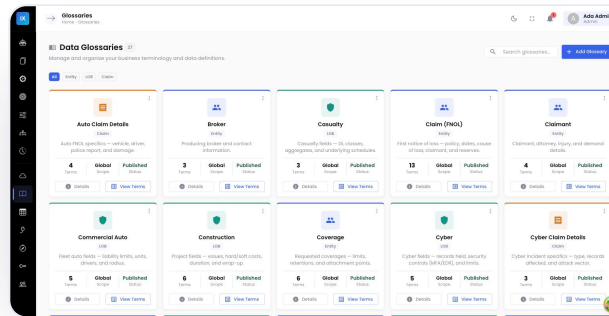
Agent overview. Many doc types → one coded output contract.



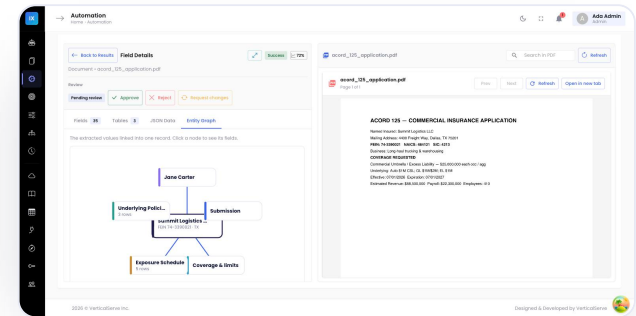
Extraction detail. Every field beside its source document — approve or reject.



Agentic pipelines. Visual designer — triggers, extraction, and delivery.



Glossaries. Business terminology & data definitions that ground extraction.



Entity graph. Extracted values linked into one connected record — click a node for its fields & source.



Try it on **your** documents — today

Ask us for sandbox access at insightxtract.com/console and extract a live submission in minutes — or download InsightXtract and install it in your own environment. Free to download and use; pay only when you want our support.

Free

to download & use

\$50K/yr

support & success (optional)

[Get sandbox access](#)

[Download & install](#)

InsightXtract — a VerticalServe product · insightxtract.com/console · Executive brief for Specialty P&C.