

Crop Urbanis Platform Technical Architecture

Source-linked records, reproducible calculations and reviewed decisions

Canonical status version	status-2026-08-26
Publication	26 August 2026

A technical path from source document to reviewed CEA decision.

The Crop Urbanis Platform is organised around a versioned project model rather than a general-purpose AI assistant. Source material, extracted claims, requirements, assumptions, calculation models, reviews and decisions remain distinct but connected. AI assists bounded information tasks; rules and calculation services apply explicit transformations; responsible people approve consequential use.

The architecture preserves distinctions that affect trust.

A source claim is not automatically a project fact

A supplier statement, client-provided value or literature reference remains a source claim until its applicability, operating conditions, evidence status and review state are recorded.

A reproducible calculation is not a guaranteed outcome

The same approved inputs and calculation-model version should produce the same result. Biological, commercial and operating uncertainty remains explicit rather than disappearing inside the calculation.

AI assistance is not professional approval

AI may extract, map, retrieve, compare and draft. Named agronomic, engineering, supplier, commercial and project authorities remain responsible for confirmation and approval within their scope.

Six functional layers

The diagram describes functional responsibilities. It does not imply that every responsibility is deployed as a separate service or that every later integration is already available.

Intake and permission boundary

Create the project and organisation context; identify authorised users and review roles; register permitted source files, versions, ownership, access conditions and intended use.

Provenance-preserving extraction

Classify documents and extract candidate values, statements and tables while retaining the original source, source location, wording, raw value, document version and extraction state.

CEA project model

Connect source claims to project objects, requirements, constraints, assumptions, alternatives, risks, owners and dependencies.

Rules and design-case services

Apply unit conversions, completeness checks, compatibility rules, physical and production calculations, utility models, financial models, sensitivity analysis and explicit project gates.

Review, change and approval

Record corrections, clarification requests, unresolved issues, professional reviews, approval states, decision rationale and the work affected by a changed input.

Output and integration boundary

Produce source-linked comparisons, assumption registers, design-case summaries, decision records and downstream project requirements; expose structured exports or integrations only where supported.

The data model links evidence, CEA objects and decisions.

The project model is not only a file index. It represents the objects and relationships needed to explain where a material value came from, how it is used, what depends on it and who reviewed the resulting decision.

Confidence and status

Extraction confidence

Extraction confidence describes the reliability of document recognition and field extraction.

Evidence status

Evidence status distinguishes measurements, supplier declarations, literature, calculations, expert estimates and assumptions.

Review status

Review status records whether the value remains unreviewed, has been corrected, has been accepted for a defined scope or remains unresolved.

Uncertainty

Uncertainty describes ranges, variability, disputed inputs and limits on prediction even where the source was read correctly and professionally reviewed.

AI works on bounded information tasks.

Language models are used where project information is unstructured or terminology varies. Each task should have a defined input scope, expected structure, validation route and review condition.

AI may assist with

- Document classification
- Candidate field and claim extraction
- Technical terminology mapping
- Retrieval from permitted project or reference material
- Draft comparison notes
- Draft clarification questions
- Multilingual transformation where separately validated

Each task is bounded by

- A named task type
- Permitted input sources
- Expected output schema
- Required source references
- Model and task version
- Validation rules
- Retry or exception limit
- Required review state

AI does not own

- Source authority
- Final unit or calculation semantics
- Confirmation of supplier warranties or compliance
- Agronomic or engineering approval
- Commercial selection
- Project acceptance
- Live greenhouse control

When a required input is missing or an output fails validation, the workflow records an exception or returns the item for review instead of fabricating completion.

The system may flag candidate inconsistencies, missing conditions and requirement conflicts. A reviewer determines whether a flagged difference is technically material.

Evaluate the components separately before expanding the workflow.

An end-to-end result can hide whether an error came from extraction, terminology, source selection, calculation, project context or professional interpretation. Evaluation is therefore divided by component and deployment condition.

Document extraction and normalisation

- Critical-field precision and recall
- Source-location accuracy
- Table and unit-normalisation error
- Missing-field recall
- Reviewer correction rate
- Reviewer time per document

Retrieval and AI assistance

- Source and citation accuracy
- Relevant-context coverage
- Unsupported-output rate
- Appropriate exception or abstention rate
- Reviewer acceptance, correction and rejection
- Performance by document and equipment class

Data and calculation services

- Required-field and schema coverage
- Dimensional and unit consistency
- Reproducibility
- Benchmark-case deviation
- Sensitivity-ranking stability
- Affected-record detection after a changed input

End-to-end workflow

- Time to a reviewable comparison
- Number of unresolved critical inputs
- Professional review time
- Decision-record completeness
- Clarification-cycle count
- Critical error escape rate

These are evaluation classes, not published performance results. Any reported result should identify the dataset or project class, baseline, metric definition, model and workflow version, review conditions and known limitations.

Internal workflow first; organisational hardening next.

The logical architecture and organisational deployment stage are not the same thing. The core workflow can operate internally while identity, isolation, monitoring, recovery and support controls continue to be hardened for selected organisations.

Current internal workflow

- Versioned project records
- Permissioned source files
- Document-processing workflow
- Retrieval and AI task services
- Rules and calculation services
- Review and decision records
- Structured output generation

Organisational hardening

- Authentication and organisation controls
- Role and permission configuration
- Tenant and workload separation
- Monitoring and audit visibility
- Data-lifecycle controls
- Backups and recovery
- Deployment automation
- Cost and operational controls

Later integration work

- Standard project-system integrations
- Operating-data feedback
- Image and sensor workflows
- Advanced analytics
- Cross-project learning under governed data conditions

Deployment planning prioritises explicit organisation, user, purpose, data, provider, region, retention, export, deletion, support and incident conditions. Detailed controls are described on the Security & Data page and agreed separately for any private-beta scope.

Cloud and programme review supplement

Condensed architecture, deployment-workload and programme-review context for technical partners.