

StructuredLayer

Construction & Property Data Readiness

CONFIDENTIAL | PREPARED EXCLUSIVELY FOR BIDSCOPE 360

Data Layer & Connected Workflows Blueprint

From scattered sources to governed workflows and AI

An implementation-ready plan for authorised portal acquisition,
continuous data structuring, one operations dashboard,
and connected workflows with controlled AI in a client-owned environment.

CONNECTED AI + BUSINESS OUTPUTS + HUMAN APPROVAL

AIRTABLE / DATABASE + OPERATIONS DASHBOARD

VALIDATION + MATCHING + DOCUMENT INTELLIGENCE

AUTHORISED PORTALS + EMAIL + APPROVED SOURCES

EXECUTIVE RESPONSE

The vision BidScope 360 shared with us

YOUR OBJECTIVE BidScope 360 wants to become the first team to see relevant construction bidding information, regardless of whether it appears in a portal, inbox, newsletter, LinkedIn source or another authorised channel. The information should already be collected, verified and organised when the team begins work - including after nights and weekends.

Document field	Current proposed value
Business requirement	Continuously identify and structure construction opportunities from five portals and additional authorised sources
Executive outcome	One dashboard showing current opportunities while background services collect, structure, verify and prepare connected workflows
Blueprint fee	\$3,500 USD fixed
Recommended implementation	\$18,000 USD before Blueprint credit
Blueprint credit	50% of \$3,500 = \$1,750 credited after implementation approval
Net implementation	\$16,250 USD plus approved third-party and usage costs
Delivery estimate	4-8 weeks after access, decisions and representative samples are available
Operating-cost scenario	Plan around \$5/day initially; optimize toward \$2/day where measurements support it

OUR DESIGN RESPONSE We recommend one connected operating layer: source evidence at the bottom, durable IDs and relationships above it, permissioned business workflows above those records, and controlled AI capabilities at the top. The first workflow proves the foundation; future workflows reuse it rather than creating new disconnected databases.

DISCOVERY BACKGROUND

What your three stakeholders told us

In our face-to-face working session, three BidScope 360 stakeholders demonstrated the workflow from the moment a source is opened to the moment a usable opportunity reaches the current table. We followed the actual clicks, waits, page changes, downloads, copy-paste work and judgement calls.

Stakeholder	What they demonstrated	Design consequence
Opportunity / bid coordinator	Portal login, saved filters, result pages, new-listing checks, file downloads and duplicate recognition	Source connector, run evidence, change detection and exception queue
Estimating / commercial lead	Which fields make an opportunity useful, how deadlines and participants are checked, and when work should move forward	Required-field gate, qualification rules and human approval
Operations / management lead	Current Airtable or spreadsheet view, morning priorities, reporting needs and the desired single dashboard	Operating database, dashboard cards, alerts, ownership and cost visibility

What we observed

- Staff visit several pages, wait for large result sets, compare new and repeated listings, open details and download supporting files.
- Some team members copy portal text or documents into ChatGPT to extract or summarise information, then copy the result back into a table.
- The same opportunity can appear again after its date, addenda, participants or documents change.
- The client wants the process to continue overnight and on weekends so the morning dashboard is already current.

DISCOVERY RECORD

The sessions behind this Blueprint

Evidence item	Recording	Retention
Opportunity and bid-coordination session	Link 1	Covered by the 14-day Blueprint decision window
Estimating and commercial session	Link 2	Covered by the 14-day Blueprint decision window
Operations and management session	Link 3	Covered by the 14-day Blueprint decision window
Approved source samples	Client-approved secure workspace	Covered by the 14-day Blueprint decision window
Transcript and workflow notes	Restricted project workspace	Covered by the 14-day Blueprint decision window

CONFIDENTIALITY

Recordings and working material are restricted to the authorised BidScope 360 and StructuredLayer project team under NDA. They are not shared, published, used as a demonstration, or used by StructuredLayer to train AI models.

CONFIDENTIALITY AND DELETION

What StructuredLayer retains after this Blueprint

Information	If implementation is approved	If implementation is not approved within 14 days
Blueprint and working design	Moves into the controlled implementation workspace.	Deleted from StructuredLayer working storage.
Call recordings and transcripts	Retained only where needed for the approved build, then deleted under the project schedule.	Deleted.
Source samples and downloaded files	Moved into the client-owned implementation environment or deleted.	Deleted.
DOM captures, console and network diagnostics	Retained only as approved technical evidence, with sensitive values redacted.	Deleted.
Contact information	Used for project communication and authorised updates.	Email address and business contact details may be retained for future StructuredLayer updates.

YOUR CONTROL


BidScope 360 may opt out of future updates at any time. StructuredLayer will then remove the retained marketing contact information, subject only to information required for legal, accounting or contractual records.

StructuredLayer does not train models on BidScope 360 recordings, source records, documents, prompts or outputs. If BidScope 360 approves a third-party model provider, only the agreed minimum data is sent under that provider selection and its approved enterprise settings.

TARGET OPERATING MODEL

How BidScope 360 will work when the layer is complete

Every three hours, a client-owned service checks authorised sources, identifies new or changed opportunities, retains evidence, structures the records, raises exceptions and publishes approved outputs to one operating view.



Outcome	What the operator sees
Morning queue	New, changed, expiring and incomplete opportunities grouped by source and priority
Opportunity record	Project, location, due dates, organisations, contacts, scope, files, source links and confidence
Document workspace	Downloaded PDFs and support files with versions, extraction status and original evidence
Exception queue	Conflicting dates, missing files, uncertain matches, failed sources and low-confidence fields
Agent status	Last successful run, next run, current task, cost, failures and approvals waiting
Downstream work	Qualified records can trigger enrichment, research, draft outreach, alerts or CRM updates

LAYERED DESIGN

Structure first; intelligence compounds above it

BidScope 360 is not buying one isolated automation. The first opportunity workflow establishes a reusable operating layer. Every later workflow can use the same project, company, contact, document, source and permission records.



Layer	What BidScope 360 owns	What becomes possible
Source and evidence	Original links, files, snapshots, timestamps and run history	Re-check facts and prove where every value came from
Structured data	Projects, opportunities, organisations, contacts, documents and change events	Search, compare, filter, report and reuse information
Identity	Durable IDs, aliases, relationships and duplicate decisions	One project or company remains one connected entity across sources
Workflow	States, owners, release conditions, permissions, exceptions and approvals	Teams and services know what can happen next
AI and outputs	Research, extraction, drafts, dashboards and approved actions	Interconnected workflows without repeatedly sending raw folders to a model

THE STRATEGIC VALUE

As the layer grows, a newly captured opportunity can update the dashboard, connect to known companies, trigger research, prepare outreach, produce management reporting and later support branded content - without rebuilding context for every workflow.

DATA LAYER SPECIFICATION

The database is the enduring product

The acquisition service, dashboard and AI models may change. The client-owned records, identifiers, relationships, evidence and event history must remain understandable and usable. This schema is designed as a durable business asset, not as the private memory of one automation.

Design principle	Blueprint requirement	Why it survives a vendor change
Client ownership	Database, files, credentials, schema and documentation remain in the client environment	A replacement team starts from an owned system, not an export request
Canonical identity	Every core entity receives an immutable StructuredLayer ID	Names, URLs and source-specific IDs may change without breaking joins
Source provenance	Every accepted value points to a source record, file, extraction and validation decision	Facts can be re-checked without trusting an undocumented model output
Event history	Important state changes are appended as events rather than silently overwritten	New workflows can understand what happened, when and why
Explicit contracts	Tables, fields, types, required rules, enumerations and triggers are documented	Another developer can integrate without reverse-engineering the dashboard
Portable interfaces	Documented API, scheduled exports and database access are available as approved	The operating interface can change without losing business meaning

CLIENT CHOICE OF PLATFORM

Airtable can provide the operational interface while PostgreSQL provides the canonical relational store at higher scale. A lower-volume first release may use Airtable as the canonical store if backups, stable IDs, field definitions and export tests are included. BidScope 360 approves the final platform; the logical schema below remains portable.

CANONICAL MODEL

Thirty tables organised into six permanent domains

This is the complete target model. The approved implementation can begin with the required subset and activate later tables without changing the core identities.


Domain	Permanent tables	Business purpose
Identity	projects, opportunities, organisations, contacts, locations	The definitive business objects used across every workflow
Acquisition	sources, source_accounts, source_runs, source_records, raw_evidence	Where information came from, how it was accessed and what each run observed
Documents	documents, document_versions, extracted_values, field_evidence	Files, versions, extracted candidates, citations and accepted values
Workflow	workflow_instances, workflow_events, assignments, approvals, exceptions	State, ownership, decisions, deadlines, failures and handoffs
Actions	agent_tasks, notifications, communications, research_findings, quote_requests, quotes	Controlled downstream work with status, cost, output and approval
Governance	project_participants, permissions, configuration_versions, cost_events, audit_log	Relationships, access, policy versions, spend and traceability

Activation plan

Release	Tables activated	Reason
Foundation	Identity + acquisition + documents + workflow + governance	Makes every captured opportunity traceable, usable and safe
Connected research	project_participants + research_findings + agent_tasks	Supports company, contractor and decision-maker research
Commercial workflows	communications + quote_requests + quotes	Supports approved outreach, quote collection and comparisons
Extended operations	Additional domain tables through controlled schema versions	Adds CRM, content, reporting or project-delivery workflows without duplicating identity

RELATIONSHIP MODEL

How the permanent records connect



Relationship	Cardinality	Meaning
source -> source_runs	one-to-many	One authorised source is checked repeatedly over time
source_run -> source_records	one-to-many	A run observes many listings and their source-native states
opportunity <-> source_records	many-to-many	One opportunity may appear in several sources or versions
project -> opportunities	one-to-many	A project may create several bid packages, notices or commercial opportunities
project <-> organisations	many-to-many	Owners, bidders, contractors, engineers and consultants have explicit roles
document -> document_versions	one-to-many	A file can be revised while retaining its identity and prior evidence
workflow_instance -> workflow_events	one-to-many	Every state change and decision forms an auditable timeline

NO NAMES AS JOINS

"ABC Construction", a portal URL or a filename is never the permanent relationship key. Joins use immutable IDs; names and source identifiers remain searchable attributes and aliases.

IDENTITY CONTRACT

The IDs that keep every workflow connected

Identifier	Example	Creation rule	Never changes when
project_id	PRJ_01K2...	Created after a project match is accepted	Project name, phase or address formatting changes
opportunity_id	OPP_01K2...	Created for a distinct commercial opportunity	Deadline, status or source listing changes
organisation_id	ORG_01K2...	Created after identity resolution	Trading name, domain or address changes
contact_id	CON_01K2...	Created for a verified person record	Title, employer or email changes
document_id	DOC_01K2...	Created for a logical document	A new revision or downloaded filename appears
source_record_id	SRCREC_01K2...	Created for one native listing identity	The same listing is observed in a later run
workflow_id	WF_01K2...	Created when a governed process starts	The workflow moves through stages
agent_task_id	TASK_01K2...	Created for one bounded machine task	It retries or changes model provider

Required identity controls

- Use UUIDv7 or ULID-style values generated by the client-owned service; never recycle an ID.
- Store portal IDs, URLs, email message IDs, domains and external CRM IDs in alias fields or mapping tables.
- Record merge decisions with surviving_id, retired_id, reason, evidence, approver and timestamp.
- Use soft deletion and lifecycle status for business records; do not silently erase a referenced entity.
- Generate folder names and dashboard links from stable IDs so documents remain connected after renaming.

CORE TABLE DICTIONARY

Projects, opportunities and locations

Table	Purpose	Primary key	Important fields	Written by
projects	Canonical identity for the physical or commercial project	project_id	name, address, coordinates, owner, sector, phase, status	Matching service; human merge approval
opportunities	One bid, package, notice or pursuit connected to a project	opportunity_id	title, scope, due_at, stage, value_band, readiness, released_at	Validation workflow; approved users
locations	Normalised project and service geography	location_id	address_raw, address_normalised, lat, lng, region, radius	Geocoder plus review
project_participants	Role-specific link between projects and organisations	participant_id	project_id, organisation_id, role, source, confidence, valid dates	Research and validation workflow

Opportunity minimum release gate

Required before release	Condition
Identity	opportunity_id and project_id exist; duplicate state is resolved
Source	At least one active source_record and original URL or message reference exists
Timing	due_at is present or explicitly marked unknown with an exception
Scope	Enough scope/category information exists for downstream qualification
Evidence	Material extracted fields point to a document, page/text region or source field
Control	readiness_status = released and the release event records actor, rule version and timestamp

BUSINESS MEANING

A project answers "what is being built?" An opportunity answers "what can we pursue?" Keeping them separate prevents one project with several packages, dates or bidders from becoming duplicate project records.

CORE TABLE DICTIONARY

Organisations, contacts and research

Table	Purpose	Primary key	Important fields	Written by
organisations	Canonical company, public body or practice record	organisation_id	legal_name, trading_name, domains, phones, addresses, status	Identity service; human merge approval
contacts	Verified person connected to an organisation	contact_id	name, title, email, phone, LinkedIn URL, verification status	Research service; approved users
research_findings	Time-bounded facts with evidence and confidence	finding_id	subject_type/id, claim, source_url, observed_at, confidence, expires_at	Research task; reviewer
project_participants	Company role on a project or opportunity	participant_id	role, source, first_seen, last_seen, valid_from, valid_to	Matching workflow

Identity resolution sequence



Rule	Implementation consequence
A found email is not automatically a trusted contact	Store verification method, verified_at and deliverability result separately
A Google Maps listing may be stale	Use it as evidence, then cross-check official website, activity and approved sources
People change companies	Employment is time-bounded; contact identity remains separate from organisation membership
Research expires	Every finding carries observed_at and optional expires_at; stale findings are not silently reused
Uncertain matches stay visible	Create a review candidate rather than forcing a wrong merge

ACQUISITION DICTIONARY

Sources, runs, records and raw evidence

Table	Purpose	Key fields
sources	One configured portal, mailbox, newsletter, page or approved channel	source_id, type, owner, access_method, terms_status, schedule, active
source_accounts	Reference to client-owned credential or session configuration	account_id, source_id, secret_reference, environment, approved_by, expires_at
source_runs	One scheduled or manual acquisition attempt	run_id, source_id, started_at, ended_at, status, records_seen, cost, error_class
source_records	The source-native listing observed during a run	source_record_id, native_id, canonical_url, first_seen, last_seen, fingerprint, change_type
raw_evidence	Minimum retained evidence needed to prove the observation	evidence_id, source_record_id, type, storage_uri, checksum, captured_at, redaction_state

A failed run is data

Run outcome	Dashboard interpretation	Allowed downstream action
success	Source completed with validation checks	Publish new and changed candidates
partial	Some pages or files failed	Publish verified records; queue exceptions; never report zero
blocked	Authentication, policy or source control prevented access	Pause source; notify; preserve last successful state
no_change	Run completed and fingerprints matched	Record heartbeat only
failed	No trustworthy completion evidence	Do not infer no opportunities; create incident and alert

ACCESS
BOUNDARY

source_accounts stores a reference to an approved secret, not the password itself. Secrets remain in a client-controlled vault or operating environment.

DOCUMENT DICTIONARY

Files, versions, extracted values and citations

Table	Purpose	Key fields
documents	Logical identity for a contract, addendum, plan, notice or support file	document_id, project_id, opportunity_id, class, title, current_version_id
document_versions	Immutable downloaded or generated file version	version_id, document_id, source_record_id, storage_uri, filename, checksum, mime_type, created_at
extracted_values	Candidate values produced by parser, rule or model	extraction_id, version_id, field_name, candidate_value, confidence, extractor_version
field_evidence	Exact support for an accepted business value	evidence_link_id, entity_type/id, field_name, version_id, page, text_region, source_url

File storage pattern

Storage layer	Stores	Example
Database	Metadata, IDs, checksums, version state, permissions and links	DOC_... current_version_id = VER_...
Client file store	Original PDF, drawings, spreadsheets and generated outputs	/projects/PRJ_.../documents/DOC_.../VER_... pdf
Evidence index	Page/text citation or bounded extraction reference	due_at supported by version 3, page 2
Operational view	Human-friendly names and links generated from records	Riverside Tower / Addenda / Current

FILES ARE NOT THE DATABASE


A folder can hold the original evidence, but it cannot reliably express identity, status, ownership, relationships, approval or change history. The database stores that meaning and points back to the file.

WORKFLOW DICTIONARY

State, events, assignments, approvals and exceptions

Table	Purpose	Key fields
workflow_instances	One governed process attached to an entity	workflow_id, workflow_type, entity_type/id, current_stage, policy_version, status
workflow_events	Append-only timeline of every material transition	event_id, workflow_id, event_type, from_stage, to_stage, actor_type/id, occurred_at
assignments	Named responsibility and due date	assignment_id, workflow_id, role, assignee_id, due_at, accepted_at, completed_at
approvals	Human decision with scope and evidence	approval_id, workflow_id, action, decision, approver_id, reason, decided_at
exceptions	Unresolved condition that prevents trusted progression	exception_id, entity_type/id, class, severity, evidence, owner, status, resolved_at

Example RFQ-to-release states



Exception example	System response	Human decision
Conflicting due dates	Hold release; show both values and evidence	Select accepted date or request clarification
Possible duplicate project	Link candidate records; prevent new project creation	Merge, keep separate or defer
Missing required file	Retry source and create visible missing-file state	Waive, obtain manually or stop pursuit
Low-confidence organisation match	Do not trigger outreach	Approve match or create a new organisation

AI TASK DICTIONARY

Models are replaceable workers, not the source of truth

Table	Purpose	Key fields
agent_tasks	One bounded extraction, research, classification or drafting task	task_id, task_type, input_refs, model_route, prompt_version, status, cost, output_ref
notifications	Message requesting attention or reporting a governed event	notification_id, event_id, channel, recipient, template_version, sent_at, delivery_status
communications	Drafted, approved or sent business communication	communication_id, contact_id, opportunity_id, channel, content_version, approval_id, status
cost_events	Measured service, model, storage or API usage	cost_event_id, task_id/run_id, vendor, units, amount, occurred_at
audit_log	Security and administrative record of reads, writes and configuration changes	audit_id, actor, action, object_type/id, before_ref, after_ref, timestamp
Model route	Use	Data rule
Deterministic code	Parsing, matching rules, checksums, calculations and writes	Preferred where rules can be expressed reliably
Local / low-cost model	Classification or simple extraction on approved data	Use where quality tests pass and sensitivity permits
Balanced hosted model	Complex extraction, summarisation or research	Send only bounded fields or approved document segments
Frontier model	Exceptional reasoning task where measured benefit justifies cost	Require explicit route, logging and approved provider settings

NO HIDDEN MEMORY

An AI output becomes business data only after it is written to a defined field or result record with task ID, model route, evidence, confidence, validation state and timestamp. Chat history is not the operating database.

PROVENANCE AND QUALITY

Every important value carries its history


Field state	Meaning	Can downstream automation use it?
observed	Value appeared in a source but has not been normalised	No
candidate	Parser, rule or model proposed a typed value	Only for non-consequential preparation
validated	Automated checks passed against approved rules	Yes where policy explicitly allows
human_accepted	Authorised user accepted the value and its evidence	Yes within the approved action boundary
conflicted	Two or more plausible current values disagree	No; exception required
stale	Freshness period expired or source changed	No consequential action until refreshed
superseded	A later accepted value replaced it	History only

Minimum provenance fields

Field	Question it answers
source_id / source_record_id	Where did this fact originate?
document_version_id + page/text region	Which exact evidence supports it?
observed_at / valid_from / valid_to	When was it seen and when was it true?
extractor_version / rule_version	Which method produced or validated it?
confidence + validation_state	How certain is it and what review occurred?
accepted_by / accepted_at	Who authorised its use?

DATABASE TRIGGERS

How records safely start future workflows



Automations subscribe to governed events. They do not watch random cells or infer readiness from incomplete records.

Governed event	Eligibility rule	Possible subscriber
source_run.completed	Run evidence confirms trustworthy completion	Deduplication, metrics and source-health dashboard
source_record.changed	Material fingerprint or document set changed	Re-extraction, deadline review and addendum alert
opportunity.released	Required fields, source evidence and approval conditions pass	Company research, dashboard priority and CRM handoff
organisation.verified	Identity and website/domain checks complete	Decision-maker research and contact verification
contact.verified	Role relevance and deliverability meet policy	Draft outreach; sending remains separately controlled
quote_request.approved	Scope, recipients and approver are recorded	Send request and start response deadline
quote.received	Message/file matched to request and scanned	Extract, compare and assign review
exception.created	Severity and owner rules resolve	Slack, Discord, email or dashboard notification

TRIGGER CONTRACT

Every subscriber is idempotent: receiving the same event twice must not create duplicate emails, folders, quotes or records. Each result writes its task ID and outcome back to the database.

FUTURE WORKFLOW MAP

What this data layer can unlock after acceptance

Future pilot	Reads from	Writes back	Human gate
Company and decision-maker research	projects, opportunities, organisations, locations	participants, contacts, research_findings	Approve uncertain identities
Personalised outreach	released opportunities, verified contacts, communication history	communications, events, delivery outcomes	Approve send policy or each message
Quote request and comparison	scope, participants, documents, deadlines	quote_requests, quotes, comparisons, exceptions	Approve recipients and commercial decisions
Proposal / PDF / deck creation	project facts, approved content, brand assets and evidence	document_versions, approvals, issued output	Approve final issue
Folder and employee-link creation	project_id, document class, permissions	storage links, access events and notifications	Approve access policy
Social and video publishing	released facts, approved graphics, source links	content drafts, assets and publish events	Approve claims and publication
Management reporting	workflow events, costs, outcomes and quality states	issued report version and commentary	Approve executive issue

**NO DOWNSTREAM
WORKFLOW
BYPASSES THE
LAYER**

A future agent receives entity IDs and a bounded set of current, permissioned records. It does not re-scrape the same context, search unmanaged folders or depend on a previous chat. This reduces duplicate work, unnecessary model context and contradictory outcomes.

PERMISSIONS AND WRITES

People and services receive only the authority they need

Actor	May read	May write	May not do without approval
Acquisition service	Configured source pages and its run state	source_runs, source_records, raw_evidence	Change business status or contact people
Extraction service	Approved document versions	extracted_values and task results	Overwrite accepted values
Research service	Released entities and approved public sources	candidate organisations, contacts and findings	Merge entities or initiate outreach
Workflow service	Current state, rules and events	events, assignments and allowed transitions	Bypass required approvals
Business reviewer	Assigned records and evidence	Accept, return, reject and correct permitted fields	Change schema or access policy
Administrator	Configuration and audit views	Approved schema, role and connector configuration	Remove history outside retention policy

PROTECTED RECORDS DURING TESTING

During parallel testing, services write to a staging environment or controlled candidate fields. Staff continue normal work in the production view until acceptance criteria pass. Cutover permissions are agreed explicitly so neither a person nor an agent can unknowingly corrupt the tested source of truth.

Sensitive fields

- Credentials are referenced from a secret manager and never stored in ordinary tables.
- Personal contact data is restricted by role, purpose and retention policy.
- Model inputs are assembled from approved fields and logged by task, provider and policy version.
- Exports inherit the same permission and deletion requirements as the source records.

HANDOVER PACKAGE

What remains if BidScope 360 uses another provider

Artifact	Delivered content	Acceptance evidence
Schema and migrations	Tables, fields, types, indexes, constraints, enumerations and version history	Clean environment can be recreated
Data dictionary	Business definition, owner, source, required rule, sensitivity and retention for each field	Business owner signs off definitions
Relationship model	Primary keys, foreign keys, cardinalities, merge rules and sample queries	Integrity tests pass
Event catalogue	Event names, payload fields, eligibility rules, subscribers and idempotency keys	Test event produces one expected result
Source catalogue	Access method, schedule, owner, validation test, failure state and recovery runbook	Each source passes supervised run
API and export contract	Documented endpoints or views plus CSV/JSON export mapping	Independent export and re-import test
Files and evidence map	Storage paths, checksums, version links, permissions and deletion rules	Random sample traces to original evidence
Operations runbook	Start, stop, monitor, retry, repair, rotate secrets, restore and escalate	Client operator completes guided exercise
Test and acceptance pack	Fixtures, edge cases, reconciliation results and signed acceptance record	Agreed measures pass

PORTABLE
HANDOVER

The handover does not require StructuredLayer to interpret private naming conventions. A competent authorised provider can understand the schema, operate the connectors and build a new workflow from the delivered contracts. Ongoing maintenance is optional and separately agreed.

RETENTION AND CONTINUITY

Permanent business records versus temporary working data

Record class	Recommended treatment	Reason
Core projects, opportunities, organisations and IDs	Retain under the client business-record policy; never hard-delete casually	They preserve identity and prevent duplicate recreation
Workflow events, approvals and merge decisions	Append-only retention aligned to audit and contractual needs	They explain why a record or action changed
Document versions and accepted evidence	Retain per project, contract and legal policy	They support traceability and future retrieval
Raw diagnostics	Short, approved retention with redaction and automatic deletion	Useful for repair but may contain sensitive technical context
Candidate extractions and rejected matches	Retain for quality analysis, then archive or delete by policy	Useful for evaluation without polluting active records
Recordings and Blueprint working files	Delete after 14 days if not approved, as stated in this document	Protects discovery confidentiality

Continuity controls

Control	Minimum requirement
Backup	Automated encrypted backup, separate failure domain and documented retention
Restore	Scheduled restore test with measured recovery time and record reconciliation
Schema change	Versioned migration, impact review, rollback path and updated dictionary
Source change	Connector version, test fixture, source health alert and supervised acceptance
Model change	Evaluation against known examples before a new model route becomes active

IMPORTANT QUALIFICATION

No database remains useful forever without ownership, backups, maintenance and controlled schema changes. This Blueprint makes continuity practical by assigning those responsibilities and delivering the documentation needed to carry them forward.

INTERVIEW FINDINGS

The workflow your team showed us

The target design follows the real workflow demonstrated by your stakeholders, including manual workarounds and the points where professional judgement must remain with your people.

Interview topic	What your team demonstrated
Sources	Five bid portals, shared inboxes, selected LinkedIn pages and authorised private industry groups
Access	Portal accounts require login; where a reliable API exists we can use it, and where it does not we can use an authorised browser workflow
Current work	Staff log in, apply filters, wait for results, compare listings, open records and download files
Repetition	The same listing may reappear after documents, dates, bidders or addenda change
Output	Staff copy selected facts into Airtable or spreadsheets and move PDFs into folders
Next process	Another person researches companies and contacts, then drafts outreach
Pain	The team sees 5-7 manually processed records where an automated pipeline may handle more
Control concern	Buyer wants one view but does not want consequential messages sent without policy and approval

Required evidence before final design

- A screen-share of each source from login through listing detail and document download.
- Twenty representative records including duplicates, changes, missing fields and failed downloads.
- The current Airtable base or spreadsheet, field definitions, filters and examples of a correct record.
- The downstream handoff: what qualifies a record, who approves it and what the next agent may do.

CURRENT-STATE DIAGNOSIS

Why adding an agent directly would fail



Failure point	Why it matters to an agent	Required response
Duplicate listings	Creates duplicate projects and repeated outreach.	Native IDs, fingerprints, matching and change events
Dates change	A captured record silently becomes stale.	First-seen, last-seen and field-level history
Portal load fails	A partial page can look like no new work.	Run evidence, checks, retries and visible failure
PDFs vary	Wrong value can land in a plausible field.	Document class, page evidence, confidence and review
Names differ	One company fragments into several records.	Identity rules, domain/address evidence and aliases
Unclear handoff	Downstream agents act before readiness.	Required-field gate and explicit release event

CORE CONCLUSION

The first automation is not outreach. It is a continuously refreshed, evidence-linked opportunity record that downstream agents can trust.

SCOPE BOUNDARY

The first production workflow

Included	Excluded unless added by change control
Five named authorised bid sources	Additional portals, groups or accounts
Three-hour scheduled checks and manual run button	Continuous high-frequency scraping
New and changed listing detection	Guaranteed capture where a source blocks or withholds access
Approved PDF and support-file downloads	Circumventing paywalls, CAPTCHA, MFA or restrictions
Structured opportunity and organisation records	Autonomous commercial judgement or bid/no-bid decisions
Airtable operating view and exception queue	Unlimited custom dashboards or a public SaaS product
Downstream release event and one agent handoff	Unapproved bulk outreach or social posting
Client VPS, accounts and documentation	Hidden vendor accounts controlled by StructuredLayer

SCOPE FREEZE

After approval, new sources, fields, actions, dashboards or communication channels require written change control. Price, timing, security and acceptance impact are agreed before work begins.

SOURCE STRATEGY

Choose the least fragile authorised connection

Priority	Method	Use when	Control
1	Supported API/webhook	Reliable authorised endpoint exists.	Schema tests, rate limits, versions and source IDs
2	Export/inbox feed	CSV, email or file contains required facts.	Fingerprint, mapping, completeness and replay
3	Authenticated request	Client is authorised and the application exposes a stable request.	Written approval, terms review, encrypted session and disable path
4	Browser automation	Routine human navigation is permitted.	Semantic controls, DOM evidence, waits, retries and checkpoints
5	Privacy-first diagnostics	A control moves, disappears or no longer responds.	Redacted DOM, accessibility, console and network signals; no screenshot by default
6	Manual review	Access or ambiguity prevents safe automation.	Owner, deadline, evidence and resolution

NO BYPASS

Session cookies can support an authorised client workflow, but they are credentials. They do not create permission, override terms, or justify bypassing MFA, CAPTCHA, paywalls, rate limits or access controls.

CONNECTION CHOICES

API and browser automation are both available

The system does not force every source into one technical method. Each source uses the most reliable authorised option, and all methods publish into the same record model and dashboard.

Option	How it works	Best fit	Tradeoff
Official API	Calls the platform endpoint with a client-owned key or OAuth connection.	Stable documented API with required fields	Version, rate-limit and usage-cost monitoring
Approved authenticated request	Uses the client-authorised application session to call the same approved data route used by the interface.	Stable request where terms and client authority permit it	Sensitive session handling and source-change exposure
Browser automation	Opens the approved browser, logs in or restores an authorised session, applies filters and visits pages.	No reliable API or export is available	Slower and more sensitive to interface changes
Privacy-first browser repair	Uses redacted DOM, accessibility, console and network signals to understand a changed interface.	Buttons move, files disappear or the next action stops working	No screenshot capture in the approved default design
Email / newsletter ingestion	Monitors an approved inbox or forwarding rule and processes new messages and attachments.	Notices arrive earlier by email than portal search	Sender validation and duplicate matching required

ONE OPERATING LAYER

API records, browser-captured records, emails and approved social-source records land in the same structure with source IDs, timestamps, files and evidence. The downstream agents do not need to know which acquisition method produced the record.

DEPLOYMENT CHOICES

Run locally, on a client VPS or in the client cloud


Option	Advantages	Responsibilities / limitations	Recommendation
Dedicated local machine	Maximum physical control; processing can remain on-site; can use a client browser profile.	Power, internet, backups, patching and authorised local access must be managed. Local does not automatically mean secure.	Choose when data policy requires on-site processing and the client can support the machine.
Client-owned VPS	Available 24/7, suitable for weekends, isolated service accounts, predictable cost and remote operations.	VPS hardening, backups, secrets, monitoring and approved outbound access.	Recommended starting balance for this workflow.
Client cloud VM	Scalable, integrates with cloud identity, logs, secret managers and private networks.	Cloud configuration, region, cost controls and client governance.	Choose when the client already operates a governed cloud environment.

BROWSER OWNERSHIP

The automated browser performs the approved repetitive navigation your user performs. It runs under a dedicated client-owned account or profile, not an employee personal profile. Its bounded self-healing layer corrects tested interface failures automatically, while unfamiliar or unsafe failures stop only the affected source and request review.

PORTAL CONNECTOR

How one no-API source is automated




Step	Agent behaviour	Evidence retained
Authenticate	Use a client-owned service account or approved session.	Account ID, status, expiry and owner - never password in logs
Filter	Select signed geography, trade, date or project criteria.	Filter version and redacted DOM or response fingerprint
Verify	Wait for expected result markers and minimum page behaviour.	Load time, page state and partial-load checks
Compare	Match portal ID and fingerprint to last successful snapshot.	New, unchanged, changed or removed classification
Retrieve	Open changed records and download approved files.	Source URL, time, checksum and original file
Publish	Write a candidate only after required validation.	Provenance, confidence, exception and run ID

FAILURE RULE

A source check succeeds only when interface, filters, results and record validations complete. A failed check is never reported as zero new opportunities.

SELF-HEALING BROWSER LAYER

Recover automatically before asking a person



If validated recovery still fails: stop that source safely, preserve state, and notify the approved channels.

Failure	Automatic recovery behaviour	Safe completion rule
Button or field moved	Re-resolve the control using role, label, surrounding DOM and expected page state.	Continue only when the intended control and resulting state are verified.
File is missing or download fails	Re-open the record, refresh metadata, retry the approved download route and check alternate approved file links.	Record missing-file exception if the source still does not provide it.
Next action does not work	Wait for pending requests, refresh, restore the last checkpoint and replay the step.	Proceed only after expected URL, record or page marker appears.
Layout or selector changed	Inspect redacted DOM, accessibility tree, browser console and approved network metadata; the repair service proposes a new route.	New route is validated against the expected business outcome before it is saved.
Session becomes invalid	Pause the source, attempt approved session renewal where configured and preserve the work queue.	Never loop credentials or bypass MFA; request authorised re-authentication when needed.
Unknown or repeated failure	Stop only the affected source, retain the last successful state and create a technical incident.	Notify Slack, Discord, email and/or dashboard according to the approved escalation policy.

PRIVACY BY DEFAULT

Recovery uses redacted DOM structure, accessibility information, console errors and approved network metadata instead of screenshots. Diagnostic payloads remove secrets and unnecessary page content. Screenshot capture is disabled in this design.

RECOVERY GOVERNANCE

Self-healing is controlled, tested and auditable

Control	How it protects continuous operation
Checkpointing	Every source run knows the last completed listing, file and database write, so repair resumes rather than starting blindly.
Bounded attempts	Wait, refresh, re-locate and AI-assisted repair each have attempt and time limits.
Expected-state validation	A repair is accepted only when the source, record and resulting page state match the signed test.
Repair history	Old selector or route, new route, reason, evidence, time and pipeline version remain visible.
Automatic rollback	A repair that fails validation is discarded and the previous stable route is retained.
Source isolation	One broken portal does not stop email ingestion, other portals or completed downstream work.
Escalation package	The notification includes source, run ID, last success, failed stage, attempts, redacted diagnostics and recovery action.
Training	BidScope 360 administrators learn how to review, approve, reject, test and manually replace a proposed repair.



CONTINUOUS OPERATION

What happens every three hours



Every 3 hours. Failed sources stay visible and never appear as zero new opportunities.

Run state	System response	Operator visibility
Success	Expected source checks and validations complete.	Green status, counts, duration and cost
Partial	Some records or files are incomplete.	Amber status and item-level exceptions
Authentication expired	Stop source worker; do not loop credentials.	Re-authentication task for access owner
Source changed	Capture evidence and stop unsafe writes.	Red status and last successful run
Rate limited	Back off within policy and reschedule.	Next attempt and affected source
Downstream blocked	Retain record but hold the agent trigger.	Release-gate reason in dashboard

The scheduler also supports a manual run, pause per source, global emergency stop and replay from retained source evidence.

DATA MODEL

The records the agents require

Record	Required fields	Why the agent needs it
Source	ID, type, owner, permitted method, frequency, status	Knows provenance and freshness
Run	Run ID, time, version, result, counts, errors, cost	Separates no findings from source failure
Opportunity	Native ID, title, location, dates, scope, value, status	One durable case for downstream work
Organisation	Name, domain, address, role, aliases, verification	Connects owners, architects, engineers and bidders
Contact	Person, role, company, source, email status, restriction	Prevents blind or repeated outreach
Document	Type, checksum, revision, dates, state, storage link	Supports grounded extraction and change review
Field evidence	Record, field, value, source, page, confidence, reviewer	Lets humans verify model output
Exception	Type, severity, owner, action, due date, resolution	Keeps uncertainty visible
Agent task	Input, allowed action, state, tool, cost, approval	Controls downstream automation

OPPORTUNITY RECORD

Representative Airtable fields and release gates

Field group	Representative fields
Identity	workflow_record_id, source_id, portal_listing_id, canonical_project_id
Discovery	first_seen_at, last_seen_at, last_changed_at, last_successful_check_at
Project	title, address, city, state, postal code, latitude, longitude, type
Commercial	value, trade scopes, bid status, due_at, timezone, addenda count
Participants	owner, architect, engineer, GC, current bidders, prior bidders, contacts
Documents	files, current issue, download status, checksum, extraction status
Quality	completeness, match status, extraction confidence, exceptions
Operations	record owner, next action, review due, qualification, release
Provenance	source URL, raw snapshot, run ID, field evidence and review history

RELEASE GATE

A downstream agent starts only when required identity, location, due date, scope and source evidence are present; blocking conflicts are resolved; and the action is allowed by the current approval policy.

TEAM ADOPTION

How people and automated services share the same layer

The operating database is not an uncontrolled spreadsheet. During testing, automated services write to protected ingestion and candidate fields while most users receive read access. Corrections and overrides enter through governed controls so the system can learn from them without losing the original evidence.

Role	Read	Write / action
Automated source service	Assigned source configuration and prior run state	Candidate records, source evidence, files, run history and detected changes
Validation service	Candidate values, matching evidence and rules	Confidence, duplicate suggestions, validation status and exception tasks
Bid coordinator	All operational opportunities and source drill-through	Qualification, assignment, review comments and approved correction requests
Estimator / commercial lead	Released opportunities, documents and research context	Business decision, scope notes, priority and approval within assigned authority
Management	Dashboard, exceptions, activity, cost and outcome reporting	Priority, approved override and escalation
System administrator	Configuration, logs, health and permission state	Approved source configuration, user access and controlled recovery

Recommended testing policy

- Raw source evidence and system-generated IDs are immutable to ordinary users.
- Users correct a value through a review action; the original candidate, correction, person, reason and time remain visible.
- Direct table editing is limited during pilot so manual changes cannot silently break matching, workflow gates or downstream agents.
- After acceptance, BidScope 360 may approve broader edit permissions for selected fields and roles without exposing protected ingestion fields.

DAILY EXPERIENCE

What changes for the BidScope 360 team

Today	After implementation
A person opens each portal and waits for results.	The source service checks approved portals on schedule and records whether each run succeeded.
Files are downloaded into local or shared folders.	Approved files are downloaded, fingerprinted, versioned and linked to the opportunity record.
Text is copied into ChatGPT and copied back into a sheet.	Extraction runs against an approved schema and returns fields with source evidence and confidence.
Duplicate or changed listings rely on memory.	Stable IDs and change events show what is new, unchanged or revised.
The team begins research after noticing an opportunity.	Released records can automatically enter company, bidder and decision-maker research queues.
Management asks for status through messages.	One dashboard shows source health, opportunities, reviews, downstream work, cost and failures.

THE OPERATING CHANGE

Your people stop being the transport layer between websites, folders, ChatGPT and spreadsheets. They become the reviewers and decision-makers supervising a continuously updated business system.

EXPECTED RESULTS

How BidScope 360 will judge whether this worked

The objective is not to claim that automation creates a guaranteed number of opportunities. The pilot establishes whether BidScope 360 can discover relevant information earlier, process more of it consistently and reduce the manual movement of data.

Result area	Acceptance measure	Evidence
Source coverage	Every approved source has a visible successful, partial or failed run state.	Run history and source-health dashboard
Freshness	Successful sources are checked every three hours, including approved weekend operation.	Scheduled run timestamps
Change awareness	New, unchanged and revised listings are distinguished without silent duplication.	Identity matches and field-level change events
Record completeness	Released opportunities contain the agreed identity, location, dates, scope, source and document fields.	Completeness score and release-gate result
Document control	Downloaded files retain source, checksum, version and current/superseded state.	Document register and evidence links
Morning readiness	The team opens one dashboard rather than beginning with five portal checks.	Daily queue and user-acceptance observation
Human effort	Routine navigation, downloading and copy-paste time are measured before and during pilot.	Baseline and pilot time sample
Operational trust	Low-confidence, conflicting and failed work appears in a review queue rather than as a completed record.	Exception and approval history
Cost control	Actual acquisition and model cost is visible by source and completed record.	Daily spend and cost-per-record view

VOLUME EXPECTATION

The new system may expose substantially more candidate opportunities than the current manual process. BidScope 360 will approve qualification rules and team capacity before additional research or outreach is allowed to scale.

IDENTITY AND CHANGE DETECTION

How repeated listings become useful updates


Check	Method	Result
Exact source identity	Source + native listing ID	Update existing source record
Document identity	Cryptographic checksum + source metadata	Retain once; link every occurrence
Project candidate	Address, title, owner, dates and location evidence	Suggest canonical project match
Organisation candidate	Domain, verified address, phone and aliases	Suggest company match
Field change	Compare candidate with last accepted value	Create before/after event with evidence
Removal or expiry	Compare with last successful source snapshot	Mark state; retain history

**NO SILENT
OVERWRITE**

Accepted values, source candidates and historical values remain distinct. When a due date or bidder list changes, the system records the event and applies the client-approved authority rule.

DOCUMENT INTELLIGENCE

How PDFs become verified fields rather than loose text




Task	Preferred execution	Cloud-model input
Checksum and duplicate detection	Local Python	None
Text extraction from digital PDF	Local parser	None
OCR for scanned pages	Approved local or cloud OCR	Only when cloud OCR is approved
Classification	Rules first; model for ambiguity	Minimum required text or page image
Field extraction	Schema-constrained with page evidence	Only approved pages or text
Date/address normalisation	Local deterministic code	None
Commercial interpretation	Named human reviewer	Optional model draft; never final

EVIDENCE REQUIREMENT

Every important value keeps the file, page, region or text span supporting it. Low confidence or conflict creates a review task instead of a guessed value.

AGENT CHAIN

What happens after a record becomes trustworthy



Agent	Allowed output	Cannot do independently
Acquisition	Listings, source files and run evidence	Bypass access or invent missing facts
Document	Class, fields, citations and confidence	Make final contractual interpretation
Validation	Completeness, duplicates, conflicts and tasks	Silently resolve uncertain identity
Enrichment	Nearby firms, roles and public verification	Call an unverified lead confirmed
Drafting	Source-linked email, CRM note, report or alert	Send externally before policy approval
Operations	Tasks, approved statuses and notifications	Approve price, contract, payment or access

EXPANSION ROADMAP

Additional agents this foundation can support

These are not included in the first \$18,000 implementation. They are the next workflows your company can evaluate once the opportunity records, documents, organisations, contacts, permissions and release gates are dependable.

Future workflow	What it can do with the structured layer	Control
Company and bidder research	Research current and past bidders, contractors, architects, engineers and suppliers connected to each project.	Retain public sources, verification state and research date
Decision-maker mapping	Find likely estimating, preconstruction, operations or executive contacts and connect them to verified organisations.	Do not represent inferred roles or unverified emails as confirmed
Contact and email validation	Validate domains, addresses and bounce status before records become outreach-ready.	Respect privacy, consent, suppression and outreach rules
Personalised outreach	Create source-grounded email and LinkedIn drafts referencing the relevant opportunity and company context.	Human approval and channel policy before sending
CRM continuation	Create or update the company, contact, opportunity, activity and follow-up date in the approved CRM.	Field-level authority and duplicate controls
Branded social publishing	Turn approved opportunities into branded posts, graphics and channel-specific captions.	Content approval, source rights and publishing schedule
Video generation	Use a client-approved avatar/video provider such as HeyGen or another API to create short branded explainers with charts and project facts.	Approved likeness, script, sources and human publishing approval
Proposal and presentation generation	Assemble approved company facts, project context, past experience and evidence into draft PDFs or decks.	Version, review and release approval
Maps and market intelligence	Show opportunity clusters, nearby firms, trade coverage, activity trends and geographic gaps.	Verified addresses, timestamped sources and clear metric definitions
Management reporting	Generate daily or weekly packs from the same governed records without manual copy-paste.	Issued version, source drill-through and named approval

FUTURE WORKFLOW PILOTS

From opportunity intelligence to commercial execution

Every record receives a stable StructuredLayer ID. That ID links the project, source, company, contact, files, research, messages, quotes, approvals and outcomes, allowing later workflows to continue without losing context.

Pilot	Automated preparation	Human / policy gate
Research remaining companies	Find organisations connected to the project, past bidders, current bidders, nearby contractors and likely delivery partners.	Verification threshold and approved source policy
Research responsible people	Map estimators, preconstruction leaders, commercial owners and other relevant decision-makers to the company ID.	Role and contact evidence review
Draft and send email	Prepare a source-grounded message, attach approved files or secure links, schedule follow-up and record the activity.	Sending can remain approval-only or move to an approved rule set
Receive RFQs and quote requests	Classify incoming requests, match project/company IDs, capture due dates and create the review task.	Commercial ownership and qualification
Request and receive quotes	Create vendor quote requests, monitor replies and attachments, connect each response to the project and package.	Recipient approval, commercial review and deadline rules
Create project folders	Generate the approved folder structure, store downloaded files separately by type/version and return the secure folder link.	Storage permission, naming and retention policy
Share with employees	Notify the assigned employee through email, Slack, Discord or dashboard with the record and folder links.	Role-based access and assignment
Prepare proposal package	Assemble approved project facts, company experience, documents, quote comparisons and draft narrative.	Estimator and commercial release

ONE ID, MANY WORKFLOWS

The project ID prevents an email, quote, downloaded drawing, bidder record, research result or employee task from becoming a disconnected object. Every downstream action writes its state back to the same connected operating layer.

COMPOUNDING VALUE

The layer becomes more useful with every connected workflow



Without the shared layer	With the shared layer
Every automation creates its own list, folders, prompts and duplicate contacts.	Every workflow reads and writes permissioned records connected by stable IDs.
The model repeatedly receives large raw files to rebuild context.	The model receives the current, relevant fields and source links required for the task.
An employee must manually move information into the next tool.	A release event starts the next approved workflow and preserves ownership.
Management sees activity after people compile a report.	Dashboard state updates as acquisition, review, research, quoting and outreach progress.

FUTURE-PILOT RULE

Future workflows are individually designed, tested and approved after the first layer is accepted. The foundation makes them faster to specify, but none is activated merely because the data exists.

EXPANSION SEQUENCE

How the second workflow starts safely



Stage	Example decision
Record complete	Opportunity has valid identity, location, due date, source evidence and current documents.
Release approved	The opportunity matches the client trade, geography, value and timing criteria.
Research	Agent searches approved public and licensed sources for relevant organisations and people.
Verify	Website, LinkedIn/company evidence, email status and restrictions are recorded.
Draft	Agent prepares the message, CRM note, social post, report or video script with source references.
Approve	A named person edits, rejects or approves the action.
Measure	Delivery, bounce, reply, engagement and outcome return to the operating record.

WHY THE ORDER MATTERS

Without a trusted project and company identity, the second agent can research the wrong business, contact the wrong person or publish an outdated opportunity. The structured release event is what makes expansion repeatable.

TOOL STACK

Client-approved, provider-neutral implementation

Layer	Recommended starting option	Alternative / decision
Environment	Ubuntu VPS with containers and encrypted storage	Client cloud account or on-premises host
Orchestration	Python workers + scheduler; optional self-hosted n8n	Approved managed workflow platform
Browser	Playwright-controlled Chromium	Supported connector or request adapter
Deterministic work	Python parsing, hashing, normalisation and validation	Equivalent client-standard runtime
AI models	Benchmark OpenAI, Anthropic and Gemini	Grok or another approved provider
Database	Airtable for operator usability	PostgreSQL for scale; Google Sheets or Excel for a bounded lightweight workflow
Files	Client SharePoint, Drive or S3-compatible storage	Bounded Airtable attachments
Interface	Client-owned Python web dashboard with queues, drill-through and run controls	Client BI, Airtable interface or internal web app
Alerts	Email, Slack, Teams or Discord connector	Dashboard-only queue
Secrets	Client-controlled secret storage and accounts	Client cloud secret manager

**NO PLATFORM
LOCK-IN**

Schema, prompts, tests, scripts, runbooks and adapters live in the client environment. A model provider can change after evaluation without rebuilding the operating record model.

AI PROVIDER DECISION

Choose models by measured task performance

Evaluation area	Representative test
Extraction accuracy	Reviewed accuracy for dates, location, organisations, contacts and scope
Citation quality	Every value points to correct page or source evidence
Vision reliability	Navigates representative layouts and stops on risky ambiguity
Latency	Completes inside the three-hour run window at expected volume
Cost	Measured per checked listing, changed record and complete opportunity
Privacy/residency	Provider terms and controls meet approved data boundary
Failure behaviour	Refuses unsupported inference and creates review item
Local-model option	A smaller local model is tested for bounded extraction or classification where data must remain on the machine
Portability	Same schema and tests run against a second provider

SELECTION RULE


StructuredLayer recommends the provider or mix that passes the agreed evaluation set. The client approves accounts, regions, budgets and data categories before production use.

MODEL ROUTING

BidScope 360 chooses the intelligence and cost profile

The core acquisition layer does not require the most expensive model for every step. Python services perform scheduling, downloads, hashing, ID generation, database writes, rules and most validation. Models are routed only to the tasks that benefit from them.

Route	Use	Typical position
High-capability frontier model	Difficult document interpretation, ambiguous repair proposals, complex research and higher-value drafting.	Highest capability and cost; selected for exceptional tasks
Balanced cloud model	Classification, structured extraction, summarisation and routine research.	Default where evaluation quality is sufficient
Low-cost cloud model	Simple field extraction, routing, tagging and repetitive bounded decisions.	Lowest cloud cost for validated simple tasks
Local small model	Privacy-sensitive classification or extraction on the client machine where benchmark quality is acceptable.	No cloud model transfer for that task; client provides suitable compute
No model	Downloads, checksums, exact matching, IDs, folders, database operations, schedules and deterministic rules.	Preferred whenever ordinary code is more reliable



YOUR SELECTION BidScope 360 may choose a premium provider, a lower-cost provider, a local model or a routed combination. Production model names and budgets are recorded in the approved configuration and can be changed without rebuilding the IDs, database or workflows.

DATA BOUNDARY

What leaves the client environment - and what does not

Processing path	Data movement	Use
Local Python	Remains on client VPS	Hashing, exact matching, rules, normalisation and routing
Local parser/OCR	Remains on client VPS	Digital text extraction or approved local OCR
Cloud model API	Selected prompt, approved text or redacted diagnostics reaches approved provider	Ambiguous classification, extraction, repair reasoning or drafting
MCP server	Receives only required arguments or resources	Expose approved database queries, files or actions
Airtable/cloud DB	Approved records stored in client account	Operations and dashboard

Client approvals required

- Which document text, redacted diagnostics, personal data and commercial fields may reach a cloud model.
- Provider retention, logging, training, region, enterprise terms and incident requirements.
- Whether any category requires local-only models, accepting possible quality or cost trade-offs.
- Redaction, minimisation, deletion, audit and human-review rules.

TRUTHFUL CLAIM

If a cloud model is used, approved input reaches that provider. “No data leaves the client environment” applies only to a genuinely local path that has been technically and contractually verified.

MCP AND CONNECTORS

How agents receive controlled tools and context

Component	Role in this design
Agent host	Coordinates model, permissions, task and approval state
MCP client	Connects the host to one approved MCP server
MCP server	Exposes narrow tools/resources such as read opportunity, list exceptions or create draft task
Portal adapter	Implements authorised API, export, request or browser method for one source
Policy layer	Checks role, record state, data category and approval before a tool runs
Audit record	Stores tool request, result, actor, source, cost, approval and failure

**WHAT MCP DOES
NOT DO**

MCP standardises how an AI application discovers tools and context. It does not create a lawful API, grant portal permission, defeat access restrictions, guarantee accuracy or replace application authorisation.

MORNING DASHBOARD

One place to supervise background work

Top-line card	Example	Action
New since last review	68	Review qualification queue
Changed records	14	Inspect dates, addenda or participants
Documents ready	51	Release eligible records
Human review required	9	Resolve confidence or matching
Failed sources	1	Open run evidence and recovery task
Drafts waiting	23	Approve, edit, reject or hold
Spend today	\$4.62	Drill into source, model and task cost

Operator controls

Control	Behaviour
Run now	Starts an approved source check and records the user
Pause source	Stops one connector without stopping the system
Emergency stop	Prevents new agent tasks and external actions
Approve/reject	Records decision, comments, evidence and actor
Replay	Reprocesses retained evidence under a selected pipeline version
Source drill-through	Opens supporting page, file, DOM evidence or raw record where permitted

SECURITY AND OWNERSHIP

Everything is built for the client to control

Control	Implementation requirement
Environment	VPS, database, storage, domains and model accounts are client-owned or transferred
Access	Dedicated service accounts, least privilege, MFA procedure and named owner
Session state	Encrypted, excluded from source control, process-restricted and revocable
Secrets	Never placed in public forms, prompts, ordinary logs or documentation
Network	Approved outbound destinations, TLS, firewall and restricted admin access
Logging	Run, tool, approval, error and configuration events without unnecessary content
Retention	Approved schedule for captures, documents, logs, model inputs and records
Handover	Code, schema, prompts, tests, runbooks, accounts and limitations
Offboarding	StructuredLayer access removed and verified

AUTHORISATION

The client confirms it may access, retrieve, process and use each source for the proposed purpose. StructuredLayer implements only the approved method and documents the disable path.

HUMAN CONTROL

The agent prepares; authorised people decide

Action	Initial production policy
Capture authorised source facts	Automatic with provenance and monitoring
Download approved documents	Automatic within source and retention policy
Suggest duplicate/company match	Automatic suggestion; human review above risk threshold
Extract dates, values and participants	Automatic candidate; review low confidence or conflict
Qualify opportunity	Rules plus human override; criteria documented
Draft email, CRM note, report or content	Automatic draft with source links
Send external communication	Human approval until separately approved policy supports narrower automation
Bid/no-bid, price, contract, payment or legal decision	Always authorised human decision
Delete records or change access	Always authorised human decision

ACCEPTANCE

The system must prove the hard cases

Test	Expected result
Known new listing	Captured once with evidence, files and required fields
Repeated listing unchanged	No duplicate; last-seen updates
Repeated listing changed	Exact before/after values and files recorded
Moved portal button	Control is re-resolved from DOM/accessibility context, validated and the run resumes
Missing file	Approved recovery routes are retried; unresolved absence becomes a visible exception
Broken next action	Wait, refresh and checkpoint replay recover the step or isolate the source
Repair cannot validate	Source stops safely and sends the approved Slack, Discord, email or dashboard alert
Partial page load	Run fails visibly; no zero-results conclusion
Expired authentication	Connector pauses and requests re-authentication
Conflicting due dates	Both values retained; review blocks release
Duplicate PDF	One object supports multiple source occurrences with provenance
Low-confidence extraction	Candidate field appears in human queue
Unauthorised agent action	Tool call blocked and audited
Provider unavailable	Bounded retry/fallback; no silent substitution
Emergency stop	No new downstream action; state remains recoverable
Cost threshold exceeded	Optional model work pauses and owner is alerted

DELIVERY PLAN

A 4-8 week implementation after approval

Week	Work	Exit decision
1	Confirm sources, terms, access, cases, fields, owners and tests.	Approved boundary
2	Build environment, schema, source register, run history and dashboard shell.	Foundation accepted
2-4	Implement and validate one portal connector at a time.	Each source passes
3-5	Build document pipeline, identity, change detection and review queues.	Records pass
4-6	Connect dashboard, release gate, one downstream workflow and alerts.	Pilot accepted
6-7	Parallel pilot; measure accuracy, coverage, latency and cost.	Readiness decision
7-8	Train team, transfer accounts/docs, remove temporary access and launch.	Signed handover

TIMING CONDITION Timing assumes agreed access and client decisions arrive inside review windows. Portal changes, unavailable permissions, added sources, migration or material security requirements use written change control.

COMMERCIAL PLAN

Implementation investment and operating cost

Commercial item	Amount / treatment
Blueprint	\$3,500 USD already paid
Recommended implementation	\$18,000 USD fixed for approved boundary
Blueprint credit	-\$1,750 USD (50%) after signed implementation approval
Net implementation fee	\$16,250 USD
Suggested milestones	50% kickoff / 30% end-to-end pilot / 20% handover
Third-party costs	VPS, Airtable, storage, model/API usage and licences paid separately
Optional maintenance	Quoted after stable handover based on sources and response level

Operating-cost scenario	Initial target	Optimised target	Assumption
Daily usage	\$5/day	\$2/day	Five sources, 3-hour schedule, changed-record processing
Approx. monthly	\$150	\$60	30-day illustration; licences and exceptional volume excluded

NOT A GUARANTEE

Daily cost is an engineering target validated during pilot. Actual cost depends on volume, page changes, diagnostics, document pages, models, retries, storage and source restrictions.

BIDSCOPE 360 COMMITMENT

What we need from your team

BidScope 360 provides	Why it is required
Full operational walkthroughs	The build must include the actual pages, waits, workarounds, exceptions and judgement used by the team.
Authorised source access	Each API, browser profile, inbox, newsletter, LinkedIn source or group requires an approved owner and method.
Representative normal and failure cases	Testing requires new, duplicate, changed, missing, restricted and broken examples.
Fast stakeholder decisions	Field definitions, qualification, permissions and exceptions cannot be completed from technical guesses.
Named business and technical owners	Access, workflow, data quality, security, acceptance and recovery each need accountable people.
Model and provider selection	BidScope 360 approves high-capability, lower-cost, local or routed model choices and budgets.
A new way of working	During pilot, staff use the dashboard, review queue and governed corrections rather than silently editing protected records.
Pilot participation	Operators compare the new system with the current process and record missed, incorrect and difficult cases.
Change control	New sources, fields, outputs and workflows are documented and approved before they enter the build.

ACCURACY PRINCIPLE

If access, workarounds, exceptions or definitions are hidden, the system will encode guesses. Quality depends on honest disclosure and authorised evidence.

ADOPTION PLAN

The end goal is dependable data without constant manual transport

Stage	BidScope 360 behaviour	StructuredLayer support
Observe	Team demonstrates the real workflow and confirms the source and output boundary.	Map steps, decisions, exceptions and access constraints.
Parallel pilot	Team continues the current process while reviewing the automated output.	Compare coverage, accuracy, freshness, failures and cost.
Controlled correction	Users submit corrections and overrides through the approved review controls.	Tune rules, IDs, extraction and recovery from recorded evidence.
Operational acceptance	Named owners sign off sources, records, dashboard, permissions and recovery.	Close accepted defects and issue the production version.
Handover	Operators and administrators run, pause, approve, repair and monitor the system.	Deliver training, runbooks, code, configuration and known limitations.
Ongoing operation	BidScope 360 owns the environment and decides whether to self-maintain or purchase support.	Optional maintenance begins only under a separately approved support scope.

CONTINUOUS DATA IS THE OUTCOME

The technology is successful only when BidScope 360 can keep receiving current, linked and reviewable information without returning to daily portal checks and copy-paste as the primary process.

HANDOVER AND MAINTENANCE

Your team receives the complete operating system

Handover item	Delivered
Architecture	Source, service, database, storage, model and dashboard diagram
Source register	Method, owner, schedule, fields, restrictions, recovery and disable steps
Data docs	Schema, dictionary, IDs, mappings, validation and authority
Agent docs	Prompts, tools, permissions, gates, models, tests and prohibited actions
Operations	Runbooks, alerts, review queue, replay, stop and cost controls
Testing	Acceptance results, known limitations and open items
Training	Operator, reviewer and technical administrator sessions
Access transfer	Client accounts confirmed; StructuredLayer access removed

Maintenance decision before handover

The implementation price includes the agreed build, testing, training, documentation and handover. Ongoing maintenance is not included. Before production handover, BidScope 360 must tell StructuredLayer whether the company will self-maintain or wants an optional support proposal. Your team will be trained to start and stop runs, review exceptions, re-authenticate accounts, approve outputs, inspect costs, review self-healing repairs and use recovery procedures. Optional support may cover unresolved source changes, connector repair, model re-evaluation, cost tuning, security updates and agreed response times.

TECHNOLOGY REFERENCES

Primary sources informing the design

Reference	Why it matters
Model Context Protocol - Architecture	Hosts, clients, servers, tools, resources and boundaries
Playwright - Authentication	Stored cookies and headers are sensitive credentials
Playwright - API testing	Browser and request contexts can share approved auth state
OpenAI API - Data controls	Cloud inputs and endpoint controls require client review
Anthropic - Computer use	Computer-use systems need execution and safeguards
Google Gemini - Computer use	Computer-use systems require controlled execution and supervision
Airtable Web API	Client-account structured record integration

Capabilities do not grant permission to access third-party services. Final design remains subject to current documentation, terms, client policy and source-specific testing.

DECISION

Approve, revise, defer, use internally or stop

Decision	What happens
Approve	Sign separate \$18,000 scope; apply \$1,750 credit; begin setup
Revise	Change sources, fields, actions, controls, timeline or commercial boundary
Defer	Retain the client-owned Blueprint for a later window
Use internally	Give the plan to the client team or another authorised vendor
Stop	End after paid Blueprint with no further obligation

Client decision record	Entry
Decision	Approve / Revise / Defer / Use internally / Stop
Conditions	
Authorised name and role	
Date	

CONTACT


Build the layer the agents can trust

StructuredLayer designs and implements client-owned operating layers for construction and property teams: authorised acquisition, connected records, human-controlled workflows, dashboards and provider-neutral AI agents.

Channel	Address
Website	structuredlayer.com
Assessment	start.structuredlayer.com
Email	start@structuredlayer.com
LinkedIn	linkedin.com/company/structuredlayer
Book a call	calendly.com/structuredlayer

Social channels

YouTube | Facebook | Instagram | TikTok | X | Threads



THE OPERATING PRINCIPLE

Get the source data right, keep it current, make uncertainty visible, then let controlled agents work on top of it.

Copyright 2026 StructuredLayer. Confidential implementation Blueprint prepared from the client discovery process.