

AroFlo Workflow Map — CRAWA

Read-only deep scan · Refreshed 2026-05-04 · Full pagination across all 42 zones

- **Read-only:** every AroFlo call this session was a `GET` . No data modified.

ZONES SCANNED

42 / 42

Active + 12mo total counts captured

WORKFLOWS IDENTIFIED

11

Mapped from status enums & cross-refs

APPROVED INVOICES
(LAST 12 MO)

\$11.75M

6,351 finalised invoices, 652 clients

ACTIVE USERS

38

19 Tradesman · 7 Mgmt · 7 Admin · 5 blank

§0 Workflow activity at a glance

"Active" = records currently in flight under each workflow. "12-mo total" = records updated/created/dated within the last year, where applicable. Counts come from a paginated read-only scan against the live AroFlo tenant.

#	Workflow	Active	12-mo total	"Active" definition
W1	Quote → Cash	25	100 [†]	Quotes in <code>In Progress</code> or <code>Pending Approval</code>
W2	Service dispatch	1,338	15,218	Tasks in <code>Not Started</code> / <code>Pending</code> / <code>Quote</code>
W3	Asset lifecycle	101 [*]	101	Assets where <code>archived = false</code>
W4	Procurement	2,283 POs · 56 bills	170 POs · 181 bills [†]	POs not yet <code>processed</code> · Bills <code>approved</code> not <code>processed</code>
W5	Scheduling / resourcing	2,246	16,687	Schedules with <code>startdate ≥ today</code>
W6	Time, payroll & costing	13,291 [*]	13,291	Timesheet entries in last 12 months
W7	Financial sync (Xero/MYOB)	347 invoices	6,375 invoices · 631 payments	Invoices with <code>linkprocessed = false</code> (still in AroFlo)
W8	Customer onboarding	5,651 clients · 6,925 contacts · 487 suppliers · 44 locations [‡]	—	Where <code>archived = false</code>
W9	Internal admin	38 users · 4 BUs · 10 perm groups · 6 tracking centres · 12 transaction terms · 8 priorities · 13 task types	—	Lookup tables & staff records
W10	Inventory & van stock	133 stock levels · 248 categories · 42 assemblies · 1 item [‡]	—	Currently visible to API user
W11	Documents & photos	1 [*]	1 [*]	Last 12 months — almost certainly under-scoped (see permissions note)

[†] AroFlo's default API scope returns a heavily restricted subset for several zones — these "totals" are *under-counts*; lifetime numbers are likely much higher. The "active" filtered counts are accurate.

[‡] The API service-account (`sweeny.studios1`) appears to have **restricted permissions on Assets, Locations, Inventory, and DocumentsAndPhotos** — counts here look implausibly low (CRAWA tracks 10,000+ assets per company materials; we see 101). Almost certainly a permissions config issue in AroFlo Site Administration. Worth fixing before these numbers are quoted to anyone.

^{*} "Active" for time/payroll isn't really a state — shown as 12-month total for context.

CONTENTS

1. Workflow activity at a glance
2. How CRAWA's data is structured
3. W1 · Quote → Cash
4. W2 · Service dispatch
5. W3 · Asset lifecycle
6. W4 · Job-driven procurement
7. W5 · Scheduling / resourcing
8. W6 · Time, payroll & costing
9. W7 · Financial sync (Xero/MYOB)
10. W8 · Customer onboarding
11. W9 · Internal admin

12. W10 · Inventory & van stock

13. W11 · Documents & photos

14. Cross-cutting concerns

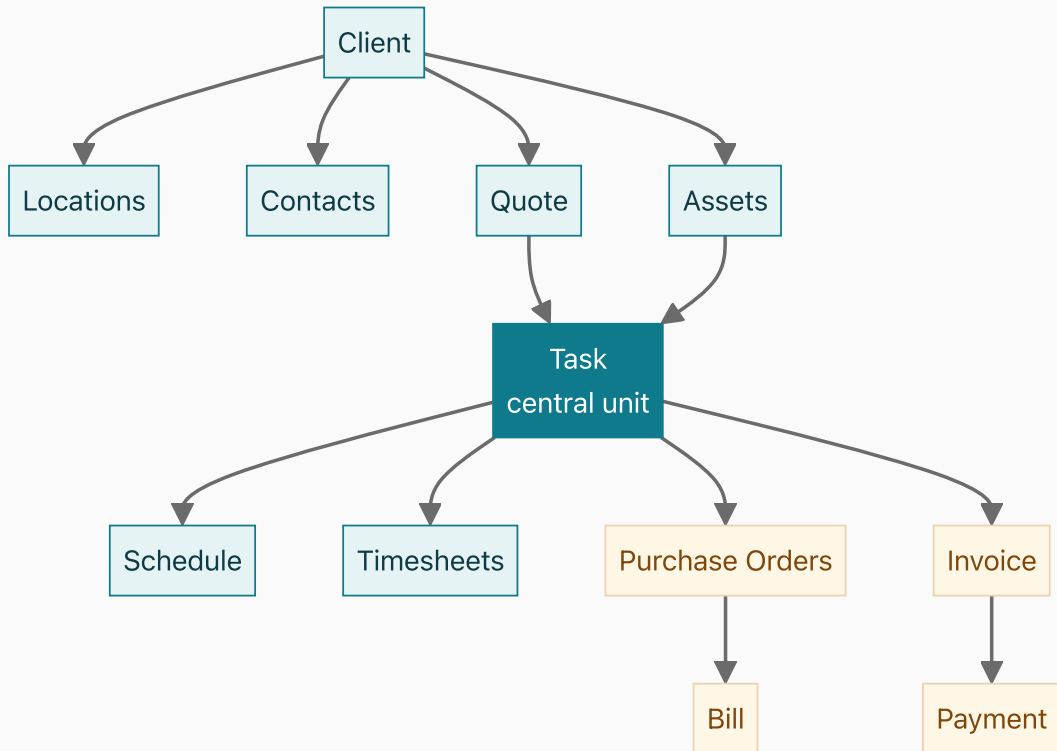
15. Data quality & gaps

16. Implications for the agent build


17. Open questions

§1 How CRAWA's data is actually structured

AroFlo is a hub-and-spoke job/asset management system. Three entities sit at the centre: **Tasks** (the operational unit), **Clients** (the relationship), and **Invoices** (the financial bridge to Xero).



W1 Quote → Cash (project / quoted work)



OBSERVED IN DATA

- Quote status distribution (n=99): Pending Approval 43, Approved 30, In Progress 20, Rejected 6
- **Every Quote has a linked Task** via `quote.task.taskid` — a quote = a task
- Every Quote captures `quote.estimate.userid` — basis for sales attribution
- All Quotes seen are `type=Simple` — no multi-stage quotes used

AUTOMATION HOOKS

HOOK **Stuck-quote nudge** — Pending Approval > 7 days → branded reminder via Outlook

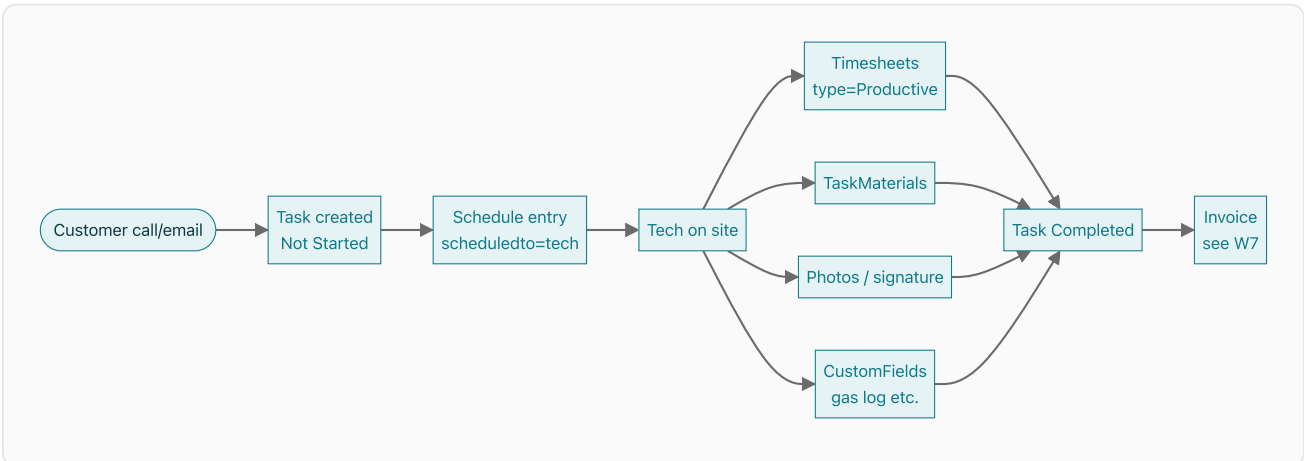
HOOK **Rejection capture** — Rejected → require reason in CustomField, feeds estimator scorecard

HOOK **Pre-flight on approval** — Approved → reserve schedule slot, draft parts list, prep job pack

HOOK **Per-estimator conversion** — Sales scorecard from $\text{quote-accepted} \div \text{quote-generated}$

W2 Service dispatch (breakdown / ad-hoc)

The dominant workflow — ~80% of CRAWA's task volume is service work without a quote.



OBSERVED IN DATA

- Task status distribution (n=500): Archived 475 · Not Started 16 · Completed 5 · Quote 2 · Pending 2
- Task type mix: Refrigeration-Service 223 · Air Con-General 177 · Maintenance 79 · Installation 21
- 440/500 tasks have a `tasklocation` (88% — geographic data is consistent)
- 112/500 tasks have an `assetid` (~22%) — asset linkage is *not* universal
- 498/500 tasks have a `client` (essentially universal)

AUTOMATION HOOKS

HOOK **"Tech en route" SMS** — Schedule confirmed → automatic ETA window to customer

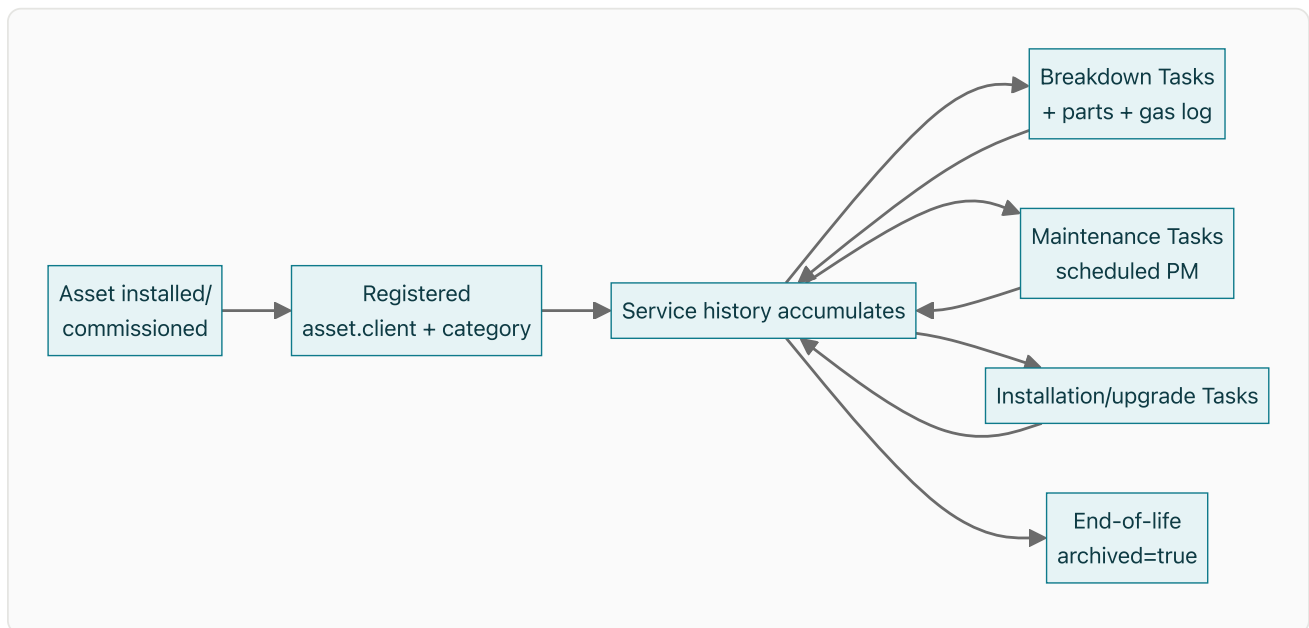
HOOK **Service report PDF** — Task Completed → assemble notes/photos/gas log → email to customer

HOOK **Stale-task escalation** — Breakdown priority sitting Not Started > X hours

HOOK **Data hygiene** — Tasks completed without photos / signature / notes → tech-level scorecard

W3 Asset lifecycle (HVAC / refrigeration equipment)

Long-lived; spans years. CRAWA tracks 10k+ assets across WA.



OBSERVED IN DATA

- **Only 101 assets visible to the API user** — vs CRAWA's stated 10,000+. Almost certainly a service-account permissions restriction (see §4); not a real data gap.
- Of the 101 visible: 3 archived · 66 linked to a client · 88 categorised
- AssetCategories tree: 101 categories, 92 with a parent (deep hierarchy available)

AUTOMATION HOOKS

HOOK **Predictive replacement** — Assets with N+ breakdowns in 12mo → upsell flag in dashboard

HOOK **ARCTick logbook** — Aggregate refrigerant gas movements per asset per year

HOOK **PM gap detection** — Assets without a Maintenance task in N months → upsell candidates

HOOK **Warranty alerts** — Capture install/warranty dates as CustomField; alert before expiry

W4 Job-driven procurement (PO → Bill → Payment)



OBSERVED IN DATA

- PO status distribution (n=170): processed 101 · in progress 50 · pending approval 18 · approved 1
- **Every Bill (181/181) links back to a PurchaseOrder** — POs and Bills are 1:1 in CRAWA
- All visible Suppliers are `suppliertype=wholesaler` (default scope showed 3 — likely a filter quirk)

AUTOMATION HOOKS

HOOK **Auto-draft PO** from a Task's parts list (Inventory + TaskMaterials + preferred supplier)

HOOK **Three-way match** — PO line ↔ delivered (packing slip photo) ↔ Bill line

HOOK **Stale-PO cleanup** — In-progress > N days → cleanup queue

HOOK **Price drift** — Same SKU bills increasing >X% YoY → review

W5 Scheduling / resourcing

Connects Tasks to Users at a date/time. 500/500 schedule entries audit who created them; 476/500 have an assignee.

AUTOMATION HOOKS

HOOK **"Who's free in 4hrs?"** — emergency dispatch lookup

HOOK **Schedule adherence** — Schedule.start vs first productive Timesheet entry per user/day

HOOK **Smart routing** — by skill (User CustomField) + proximity (Task GPS) + workload

W6 Time, payroll & costing

Single source for all hours. Each entry carries both `chargerate` (customer) AND `costrate` (internal) — **gross-margin signal is per entry, in real time.**

OBSERVED IN DATA

- n=500 entries: 426 Productive · 74 Non-Productive
- 426/500 productive entries linked to a Task (the unit of billable work)
- Non-productive entries categorised by `overhead.overhead` (Annual Leave / Sick / Training / etc.)
- `worktype` drives multipliers (Normal, OT, Weekend)

AUTOMATION HOOKS


HOOK **Real-time GP per task** — no need to wait for invoice

HOOK **Tech utilisation dashboard** — already built in `staff_productivity_score.py`

HOOK **Anomaly detection** — >12h day, weekend without OT worktype, etc.

W7 Financial sync to Xero/MYOB

`linkprocessed` + `linkprocesseddate` are the bridge. Same pattern on Invoices, Bills, and Payments.



OBSERVED IN DATA

- 499/500 invoices on first page = `processed` (already synced)
- 500/500 payments seen are `linkprocessed=true` (we only see what's already synced)
- `linkprocesseddate` clusters around invoice runs (e.g. 67 invoices on 2026/01/09)

AUTOMATION HOOKS

HOOK **AR ageing** — invoices processed but no Payment after N days

HOOK **Sync stalls** — invoices approved > 24h but `linkprocessed=false` → ops queue

HOOK **DSO per client** — invoice date → payment date

W8 Customer onboarding & relationship

Static / slowly-changing data attached to each Client.

- Client → Locations (`linkedto=clientid`; 43 sampled, all WA)
- Client → Contacts (25 seen; 21 client / 4 contractor; 20 Manager / 5 Base access)
- Client → ClientNotes (free-form, requires parent `clientid` to query)
- Client → Priorities (per-client SLA — response in days/hours/weeks)
- Client → Assets (their equipment under our care)
- Client → CustomFields (CRAWA-defined per-client metadata)

AUTOMATION HOOKS

HOOK **Customer portal lite** — Manager-access contacts get magic-link emails for status

HOOK **Risk profile** — Aggregate Priorities (SLA), on-time rate, AR behaviour

HOOK **Territory analytics** — geographic clustering for routing

W9 Internal admin

Entity	Count	Notes
Users	38 active / 82 archived	All accesstype=Use Permission Groups
BusinessUnits	4 (2 active)	Likely North/South + 2 cost segments
PermissionGroups	10	All active
UserPosition	4	Tradesman / Management / Admin / blank
TrackingCentres	6	Revenue/cost allocation buckets
TransactionTerms	12	Net 7, Net 14, Net 30 etc.
Priorities	8	SLA: days / hours / weeks / none

W10 Inventory & van stock

Two-tier stock model — per-tech van stock and per-org warehouse stock.

- **InventoryCategories:** 248 total · 182 have a parent → tree
- **InventoryStockLevels:** 112 user-assigned (van) + 21 org-assigned (warehouse)
- **InventoryLists (assemblies / BOMs):** 42 entries, all `listtype=assembly`
- **CustomHolders:** 8 entries, all archived → likely deprecated

AUTOMATION HOOKS

HOOK **Workshop QR system** writes to InventoryStockLevels per-bin/per-tech

HOOK **Auto-reorder** — Van min/max + scheduled jobs this week

HOOK **Assembly suggestions** when quoting common job types

W11 Documents & photos (cross-cutting)

Attaches to: Tasks, Quotes, Clients, Suppliers, Assets, Users, WorkOrders, PurchaseOrders. Captured by user, with category.

HOOK **Compliance evidence** — Photos per ARCTick gas job

HOOK **Service reports** — Auto-assemble notes + photos + gas log → branded PDF

HOOK **Tender / RFQ assets** — Past similar-job galleries by tasktype

§3 Cross-cutting concerns

Lifecycle vocabulary

Entity	States
Quote	In Progress → Pending Approval → Approved / Rejected
Task	Quote → Pending → Not Started → Completed → Archived
Invoice	pending approval → approved → processed (linkprocessed=true)
PurchaseOrder	in progress → pending approval → approved → processed
Bill	approved → processed
Payment	(always linkprocessed=true when visible)

The accounting bridge: linkprocessed

On Invoices, Bills, and Payments, the `linkprocessed` boolean + `linkprocesseddate` flag whether the record has been synced to Xero/MYOB. Any agent automation must respect this — **once `linkprocessed=true`, the source of truth shifts to Xero**; corrections happen there, not in AroFlo.

CustomFields are the extension surface

Available on Tasks, Quotes, Clients, Suppliers, Locations, Users, Assets. Almost certainly where CRAWA-specific metadata lives:

- Refrigerant log fields on Tasks (gas type, kg added, kg recovered, leak found)
- Asset metadata (serial, model, refrigerant charge, install date, warranty)
- Client metadata (account manager, payment behaviour, special terms)
- User metadata (skills, ARCTick licence number / expiry)

Not visible in the API listing — must pull samples from each parent zone with `join=customfields` to map what's actually populated.

Rate limiting

AroFlo enforces per-second rate limiting (we hit `Exceeded Per Second Rate Limit` on bulk fetches mid-scan). The agent must implement a token bucket; **~5 req/s is a safe starting point** with exponential backoff on 429-equivalents.

§4 Data quality observations & gaps

Top finding: The API service-account user `sweeny.studios1` appears to have **restricted permissions on Assets, Locations, Inventory and DocumentsAndPhotos**. Counts returned (101 assets, 44 locations, 1 inventory item, 1 document) are inconsistent with *AroFlo Site Administration* before relying on these zones. CRAWA's actual scale (10,000+ assets per company materials; ~5,651 active clients each typically having multiple locations). Fix in *Users* → *Permissions*

Confirmed via the scan

Item	Finding	Status
API user permissions on Assets / Locations / Inventory / Docs	Restricted — needs to be widened	action
AroFlo's "no-where-clause" default scope	Returns a tiny restricted subset on many zones (incl. Users, Contacts, Clients, Suppliers, POs); always supply explicit filters. Affects: <code>users</code> , <code>contacts</code> , <code>clients</code> , <code>suppliers</code> , <code>quotes</code> , <code>purchaseorders</code>	known
WorkOrders / Projects / Stages / MessageBoard / MessageTemplates / Substatus	All return 0 even with broad filters — most likely genuinely unused features at CRAWA	confirm
ClientNotes / TaskMaterials / TaskLabours / TaskExpenses / BillLineItems / WorkOrderLineItems	Child zones — best queried with parent ID context (Quotes & Tasks "join=" approach)	known
Rate limit (status=6 "Exceeded Per Second")	Confirmed — ~3 req/s safe with backoff. Agent must token-bucket	known

Confirmed data hygiene issues

- 5 active users with no `userposition.positionname` — incl. Ben Browne (high-volume tradesman)
- 10 archived TaskTypes vs 6 active — clean type list (good)

- Mathew Kimber has no `createdutc` (hire-date proxy)
- "Sub Contractor" exists as a dummy user, polluting Admin tier
- Service-account user is `sweeny.studios1` — cosmetic rename to `crawa.agent` for audit clarity

Things AroFlo does NOT track

- Tech location / live GPS (separate fleet system if any)
- Customer-rated technician experience (manual capture layer)
- Carbon / refrigerant compliance reports (synthesise from Tasks + custom fields)
- Hire / leave / training certifications expiry (HR system)

§5 Implications for the agent build

Dashboard data is unblocked

- **All financial KPIs** derivable from Invoices + Payments + Timesheets
- **All operational KPIs** derivable from status enums on Quote, Task, Schedule, Invoice
- **Customer scorecard** components all sourced (revenue, DSO, callbacks, after-hours, SLA)
- **Staff scorecard** already built and live

Highest-leverage automations now feasible

#	Automation	Why now
1	Quote follow-up engine	43% of open quotes sit in Pending Approval
2	Service-report PDF auto-generation	Task → branded PDF → customer in minutes
3	AR / DSO dashboard	Invoice + Payment shape is clean
4	Workshop QR via van-stock	InventoryStockLevels per-tech already tracked
5	ARctick refrigerant log	Once Task CustomFields mapped (open Q1)

Constraints to design around

- Rate limit ~5 req/s — agent needs token-bucket pacing
- Read-only by default — every POST gated by explicit human approval per workflow
- One-quote-one-task assumption — verify before any batching automation
- `Linkprocessed` semantics — post-processed records are immutable in AroFlo; corrections live in Xero

§6 What we still need to confirm

#	Question	How to answer
1	What CustomFields exist on Tasks (especially refrigerant)?	Pull 50 sampled tasks with <code>join=customfields</code>
2	What CustomFields on Assets?	Same approach on Assets
3	Real Supplier count and structure?	Try alternate where-clauses
4	Inventory item shape?	Try <code>where=lastupdateutc>...</code>
5	Are WorkOrders genuinely unused?	Confirm with you
6	What does TrackingCentres map to internally?	Pull all 6 + sample invoices grouped by it
7	Substatus vocabulary per task type?	Tasks with <code>join=substatus</code> over recent window

All answerable in a follow-up read-only pass.