

Inbound Leads, Researched and Routed Before the Prospect Closes the Tab

An AI lead-qualification system built to run unattended on real revenue · Case study

< 7 sec form submitted to team notified, across live runs	100% of leads researched, scored, routed and logged	0 duplicate alerts inside the protection window	Every hard failure reaches a person, with a link
Client	B2B AI-automation agency (name withheld)		
Engagement	Design and build an AI lead-qualification pipeline		
Systems connected	CRM · team chat · company-data enrichment · AI scoring · audit log		
Delivered	End-to-end qualification pipeline plus a reliability layer: failure alerting, duplicate blocking, and an audit trail		
Status	Running. Ideal-customer criteria and team routing each change in one place, nothing else moves		

The problem

Every inbound lead arrived looking identical: a name, an email, a short message.

Which meant the founder with an approved budget sat in the same queue as someone browsing out of curiosity. To know who deserved the first call, someone had to research each company by hand — its size, its industry, whether the message meant anything. That work takes time, and it usually got done after the lead had already gone cold.

The ask wasn't "add AI." The real requirement, which surfaced during scoping, was stricter: a system safe enough to run unattended on real revenue-bearing leads — one that never pings the team twice for the same person, never drops a lead silently, and says something the moment it breaks.

What the system does


One system takes over the moment a lead arrives.

- 1. It filters before it spends.** Personal email addresses skip the paid company-research step entirely — budget goes only toward leads worth researching.
- 2. It researches the company.** Size, industry, funding, age — none of which was in the form submission.
- 3. It scores and tiers.** Every lead gets an A/B/C tier and a 0–100 buying-intent score, each with written reasoning a salesperson can read and argue with — not an unexplained number.
- 4. It routes to the right place.** A-tier hits the urgent sales channel with a follow-up task due in 4 hours. B-tier goes to the standard channel with a 2-day task. C-tier enters long-term nurture with a 7-day task.

5. It writes the first sentence. Every notification carries an opening line written for that specific lead, so nobody starts from a blank page.

6. It logs everything. Each processed lead becomes one row in a sheet the client opens directly, without touching the technical system.

The practical result: the phone buzzes before the prospect has closed the browser tab.



The full pipeline. Core path in purple, reliability layer in green.

What it looks like in practice

Every screen below is a live run, not a mockup.

An A-tier lead: 80/100 buying intent, the reasoning behind the score, and an opening line the rep can send as-is — routed to the urgent sales channel.



The same system, a different verdict. A decision-maker at a company outside the ideal size range scores B and lands in the standard channel, not the urgent one.

Name	Email	Create Date	icp_tier	intent_score	data_confidence
Dania Mansour		Today at 4:39 PM GMT+3	B	70	medium
Karim Sabbagh		Today at 4:17 PM GMT+3	B	70	medium
Hana Osei		Today at 4:14 PM GMT+3	B	60	medium
Bilal Younes		Today at 4:10 PM GMT+3	C	0	low
Ronald J		Today at 11:16 AM GMT+3	A	0	medium
heba saad		Today at 9:57 AM GMT+3	C	0	medium
mo		Today at 5:37 AM GMT+3	B	0	medium
Sami Rahal		Jul 31, 2026 12:53 PM GMT+3	C	0	low
Tomas Berg		Jul 31, 2026 12:52 PM GMT+3	B	40	medium
Rania Haddad		Jul 31, 2026 12:50 PM GMT+3	A	70	medium
Layla Fahad		Jul 31, 2026 5:18 AM GMT+3	C	0	medium

Tier, buying-intent score and data confidence written straight onto the CRM record — no manual entry.

	C	D	E	F	G	H	I	J	K
	name	company	title	icp_tier	intent_score	data_confidence	was_enriched	slack_channel	hubspot_task_id
	Natalie Kim	Figma	VP of Sales	C	70	medium	TRUE	#marketing-nurt	508636592369
	Jordan Reyes	Buffer	Director of Oper	A	70	medium	TRUE	#sales-hot	508616006845
	Priya Nair	Typeform	Marketing Mana	B	40	medium	TRUE	#sales-warm	508638997728
	Layla Fahad	Notion	Ops Lead	C	0	medium	TRUE	#marketing-nurt	508592798945
	Tomas Berg	Mixpanel	Marketing Mana	B	40	medium	TRUE	#sales-warm	508641684691
	Sami Rahal	N/A	Freelancer	C	0	low	FALSE	#marketing-nurt	508675680484
	Rania Haddad	Linear	VP of Revenue	A	70	medium	TRUE	#sales-hot	509160668393
	mo	Broadly	vp of ops	B	0	medium	TRUE	#sales-warm	509109931223
	heba saad		partner	C	0	medium	TRUE	#marketing-nurt	509389827277
	Ronald J	WhereTo, a Fligh	VP of Ops	A	0	medium	TRUE	#sales-hot	509324717296
m	Bilal Younes	na	Freelancer	C	0	low	FALSE	#marketing-nurt	509409496252
	Hana Osei	Airtable	Marketing Mana	B	60	medium	TRUE	#sales-warm	509463145689
	Karim Sabbagh	Vercel	VP of Sales Ope	B	70	medium	TRUE	#sales-warm	509431240897
	Dania Mansour	Retool	VP of Revenue	C	70	medium	TRUE	#sales-warm	509453360349


One row per lead: identity, company, tier, score, data confidence, the channel it was routed to, and the follow-up task it created.

What makes it safe to leave running

Most systems like this stop at “lead arrives → AI scores it → someone gets a message.” This one ships with three safeguards, each tested against real failure conditions before handover.

Instant failure alerting

When something breaks, an alert names the exact step that failed, why, and links straight to the failed run. Verified by deliberately breaking the live system and watching a real alert arrive — not a simulated one.




A real alert from a deliberately broken run: the failing step, the reason, and a direct link to the execution.

Duplicate blocking

Double form submissions are a fact of life. The system blocks the same lead being processed twice within 24 hours — before it can spend research budget, interrupt the sales team a second time, or create a duplicate task. Verified by firing an identical submission twice: the first processed fully, the second stopped cold with no side effects.

A readable audit trail

Every lead becomes one row: date, identity, tier, score, data confidence, and where it was routed. The client reads it directly — and the same table is the memory the duplicate check relies on. One artifact, two jobs, nothing to keep in sync.



Run history alongside a single execution opened up: the score, the written reasoning, and the channel it chose. Every live run finished in under seven seconds.

Three decisions that changed the outcome

The AI proposes; code decides

During testing, the model was instructed not to exceed a certain score when company data was thin. It exceeded it anyway. A model costing roughly 13× more exceeded it by a wider margin. The fix wasn't paying more — it was a hard limit enforced in code, which has held on every run since.

The commercial consequence: the cheaper model became the correct choice, because the guarantee no longer depended on the model's judgment at all.

Cost control is architectural, not an afterthought

Personal email addresses never reach the paid research service. Weak matches are rejected at the source, before they are paid for. Spend tracks the value of the lead rather than the volume of traffic.

A silent failure was caught in testing, not production

During duplicate-blocking tests, a first-time lead — under one specific condition — could exit the pipeline with no error message at all. Every genuinely new lead would have vanished without a trace. It surfaced because failure paths were tested, not just success paths, and it was fixed before launch. That is the concrete value of testing what can break rather than only what should work.

How it was verified

A deliberately awkward test set: a large global company whose contact was not a decision-maker (correctly scored low despite the company's size), a personal email address, a high-urgency lead with thin company data (the score-limit stress test), and clean A- and B-tier profiles. Each one was checked in four places — the notification, the CRM record, the created task, and the audit log.

Then the failure paths: a real production failure to prove the alerting, and a two-shot submission to prove duplicate blocking.

Results

- A complete manual triage step is now automatic: every lead is researched, scored, routed, tasked and logged with no human involvement.
- Form submission to team notification in under seven seconds, measured across live runs.
- Duplicate notifications and duplicate tasks are impossible inside the protection window — by design and by test.
- Any hard failure reaches a person with a direct link to the problem.

What changes when your business changes

The two things most likely to change — the definition of an ideal customer, and who follow-up tasks are assigned to — are each isolated to a single point. When the client's criteria evolve or their team changes, nothing else in the system needs touching.

If your inbound looks like this

Built for one agency, but the pattern transfers to any business where leads arrive faster than anyone can qualify them and the good ones cool off in a shared inbox. The system is rebuilt around your CRM, your channels, and your definition of a good lead.

Mohamed Altahan · SoftecNova · mohamedaltahan.com · softecnova.com