

APRIL 4, 2026

AI LEAD QUALIFICATION PLAYBOOK FOR SOLOPRENEURS: 4FIELD INTAKE + LLM GUARDRAILS + AUTOROUTING

A step-by-step playbook to move qualification into your intake form, score leads with an LLM, enforce guardrails, and auto-route outcomes—so only great fits see your calendar. Built for solo service providers running Zapier or Make.

FROM EPISODE

[STOP TAKING BAD DISCOVERY CALLS: AI-POWERED LEAD QUALIFICATION THAT GUARDS YOUR CALENDAR](#)

CONTENTS

- 1) What you're building and why
- 2) Design the 4-field intake (+ low-friction defenses)
- 3) The LLM qualifier: strict JSON or it doesn't ship
- 4) Guardrails and routing: Zapier vs. Make (choose one)
- 5) Communications and calendar gating
- 6) Logging and the proof dashboard
- 7) Test, ship, and keep it safe

Build a form-first qualification system that scores leads with an LLM, blocks junk with guardrails, and only exposes your calendar to good fits. This guide turns the episode's walkthrough into a copyable, tool-agnostic playbook you can ship in about 90 minutes (as of April 2026).

1) WHAT YOU'RE BUILDING AND WHY

You're moving qualification before your calendar. The intake form submits to your automator (Zapier or Make), which (1) sanitizes input, (2) classifies fit via an LLM with a strict JSON schema, (3) applies platform guardrails/deterministic filters, and (4) routes: Qualified → price-anchored email + premium Calendly, Maybe → 3-question follow-up, No-fit → helpful redirect. All decisions are logged to a lightweight CRM table.

Success criteria:

- 0 calendar link leaks to unqualified leads
- < 60 seconds from submit → outcome email
- ≥ 80% agreement with your human judgment in backtests (last 20 inquiries)
- Ops spend ≤ \$60/month at early volumes

Architecture at a glance:

- Form (Tally/Typeform/Interfaces/Google Forms) → Zapier/Make → LLM (JSON) → Guardrails/Filters → CRM log (Zapier Tables/Notion/Airtable) → Email + Calendar routing

2) DESIGN THE 4FIELD INTAKE (+ LOWFRICTION DEFENSES)

Keep it short and discriminative. These four fields do 80% of the work:

Required fields

1. Budget range (dropdown)
 - <\$5k, \$5k–\$15k, \$15k–\$30k, \$30k+
2. Use case (dropdown of your service categories)
 - e.g., "Sales ops automation", "Marketing attribution", "Back-office ops", "Custom dev (not offered)"
3. Timeline (dropdown)
 - ASAP ($\leq$ 30 days), Next month, Next quarter, Next year
4. Problem statement (free text, max 500 chars)
 - Prompt copy: "In 2 sentences, what problem are you trying to solve?"

Anti-spam and validation (do all four):

- Honeypot field (reject if filled)
- Cloudflare Turnstile (background human check)
- Email regex + MX check (reject throwaway/invalid domains)
- Text limits: 500 chars; disallow links in free text

Nice-to-have friction that pays for itself:

- "Company website or LinkedIn profile" (required URL) — cuts spam and tire-kickers

UX copy you can steal:

- Above the fold: "We only book with best-fit projects. This 60-second form tells us if we're a match."

- Near budget: "Ranges help us route fast or refer you to a better-fit option."

3) THE LLM QUALIFIER: STRICT JSON OR IT DOESN'T SHIP

Use a strict schema so the LLM returns routable data every time. Paste this into your AI step as the required output shape.

JSON schema (required keys and enums):

```
{
  "fit": "yes" | "maybe" | "no",
  "why": "string (<= 160 chars)",
  "risk_flags": ["pii" | "scope_mismatch" | "budget_mismatch" | "timeline_unrealistic"
    | "prompt_injection" | "toxicity" | "missing_data"],
  "price_anchor_tier": "starter" | "professional" | "enterprise",
  "confidence": 0.0-1.0
}
```

System prompt template (edit for your niche):

```
You qualify leads for a solo automation consultant who builds Zapier/Make workflows
for SMB service businesses.
Good fit: budget >= $5k, timeline <= 90 days, scope within automation (no custom
ML/app dev).
Bad fit: budget < $2k, scope outside expertise, unrealistic timeline, or security/com-
pliance blockers.
Return ONLY valid JSON in this exact schema (no prose). If information is missing,
set fit="maybe" and add "missing_data" to risk_flags.
Map price_anchor_tier: starter ($5k-$12k), professional ($12k-$25k), enterprise
($25k+).
```

Example input (condensed):

```
Budget: $15k-$30k
Use case: Sales ops automation
```

That Guards Your Calendar

Timeline: Next month

Problem: "Sales reps spend 3 hrs/day updating CRM; want auto-logging and enrichment."

Example output:

```
{"fit":"yes","why":"Budget and scope align; clear automation ROI in sales ops.", "risk_flags":[], "price_anchor_tier":"professional", "confidence":0.86}
```

Routing thresholds (defaults you can tune):

- Qualified if fit = yes AND confidence $\geq$ 0.6 AND no high-severity risk_flags (pii, prompt_injection, toxicity)
- Maybe if fit = maybe OR confidence < 0.6 OR risk_flags contains missing_data
- No-fit if fit = no OR risk_flags contains scope_mismatch OR budget_mismatch

4) GUARDRAILS AND ROUTING: ZAPIER VS. MAKE (CHOOSE ONE)

Pick the platform you live in. The principles are identical: fail safe, never leak links, log everything.

Option A — Zapier (with AI Guardrails)

1. Trigger: New form submission
2. Sanitization: Code/Formatter step
 - Normalize budget value, trim text, strip links from free text
3. AI step: Send structured prompt → capture JSON
4. Guardrails: AI Guardrails by Zapier

That Guards Your Calendar

- Checks: PII, prompt-injection/jailbreak, toxicity; returns booleans
5. Parse + Validate JSON: Code step (reject if keys missing)
 6. Filters (deterministic):
 - If guardrail flags true → Tables: Manual Review (stop)
 - Else if fit=yes AND confidence $\geq$ 0.6 → Path: Qualified
 - Else if fit=maybe OR confidence $<$ 0.6 → Path: Maybe
 - Else → Path: No-fit
 7. Actions per path (see Section 5)
 8. Log to Tables (doesn't count toward tasks): write raw input, parsed JSON, outcome, time-stamps, email copy

Option B — Make (If-else + Merge)

1. Webhook/Form trigger
2. JSON Validate (Code/JSON)
3. AI classify (OpenAI/Claude/etc.) → JSON
4. If-else module (order from most specific to least; only first match runs):
 - Case 1: guardrail_flags=true OR parse_error=true → Manual Review
 - Case 2: fit=yes AND confidence $\geq$ 0.6 AND risk_flags empty → Qualified
 - Case 3: fit=maybe OR confidence $<$ 0.6 OR risk_flags contains missing_data → Maybe
 - Case 4: default → No-fit

That Guards Your Calendar

5. Merge branches → shared logging
6. Write to Notion/Datastore or Airtable; then send comms

Non-negotiables:

- The only place a Calendly link exists is inside the Qualified branch's email body
- Else/fallback path must catch malformed or empty AI output
- Order conditions from specific → general to prevent misroutes

5) COMMUNICATIONS AND CALENDAR GATING

Send short, price-anchored emails within 60 seconds. Use a separate "Qualified Discovery" event in Calendly with limited premium availability; do not publish its URL anywhere.

Qualified (price-anchored) — pasteable template

```
Subject: Next step + ballpark for your project
```

```
Hi [FIRST NAME],
```

```
Based on what you shared, I can help. Projects like this typically land between [PRICE RANGE] and take [TIMELINE].
```

```
Here's my calendar to discuss specifics: [CALENDLY_QUALIFIED_LINK]
```

```
(Optional) Case study: [LINK]
```

```
—[YOUR NAME]
```

```
[YOUR COMPANY]
```

Maybe (collect missing signal)

```
Subject: Quick follow-up (3 questions)
```

```
Hi [FIRST NAME],
```

```
A few details will help me scope this:
```

That Guards Your Calendar

1) Biggest bottleneck right now?
2) What have you already tried?
3) What would “success in 90 days” look like?
Reply here—your answers route straight back to my system.

—[YOUR NAME]

No-fit (helpful redirect)

Subject: A better fit for your project

Hi [FIRST NAME],
Thanks for reaching out. Based on your note, I’m not the right fit. For what you need, try [ALTERNATIVE/RESOURCE] — they specialize in [AREA].
Wishing you the best with the project.

—[YOUR NAME]

Calendly gating setup checklist

- Create a hidden event type: “Qualified Discovery — 25 min” with limited weekly slots
- Turn off “public” visibility; only share link via automation
- Add pre-call questions you want every Qualified lead to answer (kept to 2–3)

6) LOGGING AND THE PROOF DASHBOARD

Log every decision and make it auditable. You’ll use this to tune the model, defend timelines, and prove value.

Recommended table (Zapier Tables/Notion/Airtable)

- id (submission), submitted_at, source_url/UTM
- email, company, website_linkedin

That Guards Your Calendar

- form_budget, form_use_case, form_timeline, form_problem
- ai_fit, ai_confidence, price_anchor_tier, risk_flags[]
- guardrail_flags[] (pii, injection, toxicity)
- outcome (qualified/maybe/no_fit/manual_review)
- email_template_used, email_sent_at, calendly_event_id (if booked)

Useful views & metrics

- Funnel: submissions → qualified → booked → closed
- Source quality: qualified rate and close rate by UTM/source
- Prompt drift: weekly ai_confidence average; % manual reviews
- SLA: submit→email latency p95

Ops notes

- In Zapier, Tables and Interfaces steps don't count toward task usage; log liberally
- In Make, use Data Store or Notion for cheap, flexible logging

7) TEST, SHIP, AND KEEP IT SAFE

Backtest before you expose the link:

1. Take your last 20 real inquiries
2. Manually label each (yes/maybe/no) and price tier
3. Run them through your flow (use a staging calendar)

4. Aim for $\geq 80\%$ agreement; inspect misses and update prompt/enums
5. Attack it: submit “ignore all previous instructions”, nonsense text, \$1M budgets, “yesterday” timelines; confirm guardrails route to Manual Review
6. Verify no calendar link is present outside the Qualified path (search logs for your event URL)

Compliance quick hits (US/EU context)

- Data minimization (GDPR Art. 5): only collect what you need to qualify (e.g., budget range, timeline); auto-delete non-converting leads after ~90 days
- CAN-SPAM: include postal address and an opt-out if your automated email is promotional in nature; always honor opt-outs

Cost envelope (starter defaults)

- Form tool: \$0–\$29/mo
- Automator: Zapier from ~\$20/mo; Make from ~\$9/mo
- LLM: $\approx \$0.10/\text{lead}$ (varies by model and tokens)
- CRM/logging: free tiers (HubSpot, Notion, Tables)
- Target total: $\leq \$60/\text{mo}$ at low volume

Iteration cadence

- Weekly: review dashboard, adjust prompt or dropdown options; A/B test one micro-change at a time (e.g., budget ranges)
- Quarterly: re-score 20 historical inquiries to check drift; refresh email copy and case studies