

PULSE

API Documentation

Predictive deal-risk scoring and AI-powered re-engagement

Version 1.0

CONTENTS

Pulse API Reference

| | | |
|-----------|-------------------------------------|---|
| 01 | Overview..... | 3 |
| | Tech stack..... | 3 |
| 02 | Risk scoring..... | 3 |
| | Scoring rules..... | 3 |
| | Bucket thresholds..... | 4 |
| 03 | Endpoints..... | 5 |
| | List all deals..... | 5 |
| | Get a single deal..... | 5 |
| | Pipeline summary & leaderboard..... | 6 |
| | Draft AI re-engagement email..... | 6 |
| 04 | Error responses..... | 8 |
| 05 | Real-world integration..... | 8 |
| | Replacing the data source..... | 8 |
| | Supported integration targets..... | 9 |
| | Additional signals..... | 9 |
| | Authentication..... | 9 |

Overview

The Pulse API powers an AI-driven sales pipeline dashboard. It scores active deals against five transparent risk signals, exposes summary metrics and rep leaderboards, and generates contextual re-engagement emails using a hosted large language model.

All endpoints return JSON. No authentication is required in the current demo deployment.

Tech stack

| | |
|--------------------|-----------------------------------------------|
| Framework | FastAPI (Python 3.10+) |
| AI provider | Groq API — Llama 3.3 70B |
| Data source | Synthetic CRM dataset (deals.json, 200 deals) |
| Hosting | Railway |

Risk scoring

Every deal is scored 0–100 against five transparent rules. Each rule emits a point value and a human-readable reason that surfaces directly in the dashboard.

| Field | Type | Description |
|--------------------|----------|--------------------------------------------|
| customer_silence | 0–30 pts | Days since the customer last replied |
| stage_stagnation | 0–25 pts | Time in current stage vs team median |
| stale_next_step | 0–15 pts | Overdue or missing next step |
| single_threading | 0–15 pts | Number of engaged contacts at the account |
| close_date_slipped | 0–15 pts | Whether the expected close date has passed |

Bucket thresholds

| Field | Type | Description |
|--------|-------|-------------------------------------|
| green | 0–35 | Healthy — no action required |
| yellow | 36–65 | Watch closely — early warning signs |
| red | 66+ | Needs intervention this week |

Endpoints

List all deals

GET /api/deals

Returns every open deal with its score, bucket classification, and per-rule risk breakdown. Results are pre-sorted by risk score descending so the riskiest deals appear first.

Example response

```
[
  {
    "id": "DEAL-0012",
    "account_name": "Meridian Networks",
    "owner": "Nicole Brennan",
    "amount": 100000,
    "stage": "Negotiation",
    "score": 100,
    "bucket": "red",
    "breakdown": [
      { "rule": "customer_silence", "points": 30, "reason": "Customer last replied 47 days ago" },
      { "rule": "stage_stagnation", "points": 25, "reason": "In Negotiation for 63 days (3.5x team median)" }
    ]
  }
]
```

Get a single deal

GET /api/deals/{deal_id}

Returns one deal with its complete activity timeline, recent email snippets, and scoring breakdown. Returns 404 if the deal ID is not found.

Path parameters

| Field | Type | Description |
|---------|--------|---------------------------------------|
| deal_id | string | Deal identifier in the form DEAL-NNNN |

Pipeline summary & leaderboard

GET /api/summary

Returns pipeline-wide KPIs (total pipeline, pipeline at risk, deal counts per bucket, average days of customer silence) and a ranked leaderboard of reps by at-risk dollar amount.

Example response

```
{
  "total_deals": 200,
  "total_pipeline": 8865000,
  "pipeline_at_risk": 3160000,
  "red_count": 21,
  "yellow_count": 34,
  "green_count": 145,
  "avg_silence_days": 8.2,
  "rep_leaderboard": [
    { "owner": "Nicole Brennan", "at_risk_amount": 595000, "at_risk_count": 10 }
  ]
}
```

Draft AI re-engagement email

POST /api/deals/{deal_id}/draft-email

Generates a contextual re-engagement email using Llama 3.3 70B via the Groq API. The prompt is constructed from the deal's risk breakdown, recent customer/rep email history, current stage, and dollar amount.

Typical response time is under 2 seconds. Requires the GROQ_API_KEY environment variable to be set on the server.

Example response

```
{
  "subject": "Quick check-in on the Meridian proposal",
  "body": "Hi team – wanted to circle back on the proposal we sent over last
month.\n\nI know Q2 planning is in full swing on your end. If it would help, I'm
happy to set up a 15-minute call to answer any open questions or revisit pricing
based on where you've landed internally.\n\nLet me know what works best."
}
```

Error responses

Errors follow FastAPI's standard format. The HTTP status code indicates the category and the detail field describes the specific failure.

| Field | Type | Description |
|-------|--------|--------------------------------------------------------------------|
| 404 | string | Deal not found — invalid deal_id passed to a deal endpoint |
| 500 | string | Server error — typically a missing API key or upstream LLM failure |

Example

```
{
  "detail": "Deal DEAL-9999 not found"
}
```

Real-world integration

Pulse currently runs on a synthetic dataset of 200 deals generated by `data_generator.py`. The architecture is designed so this data layer can be swapped for a live CRM integration without changes to the scoring engine, API surface, or frontend.

Replacing the data source

The scoring engine in `scoring.py` operates on a normalized deal schema — account, owner, amount, stage, dates, and recent activity. Any CRM that exposes these fields can feed Pulse directly. The recommended path is to add a thin adapter module that pulls from the target CRM on a schedule and writes into the same schema the API already consumes.

Supported integration targets

| Field | Type | Description |
|------------|----------|-----------------------------------------------------------------------------------------|
| Salesforce | REST API | Pull opportunities, accounts, activities, and contact roles via the Salesforce REST API |
| HubSpot | REST API | Pull deals, companies, engagements, and associations via the HubSpot CRM API |
| Pipedrive | REST API | Pull deals, activities, and persons via the Pipedrive API |
| Custom CRM | Adapter | Implement a custom adapter that maps internal fields to the Pulse schema |

Additional signals

With real CRM data connected, the scoring engine can be extended with signals that aren't possible to simulate, such as email reply sentiment from call recording platforms, meeting frequency from connected calendars, multi-threading depth across the buyer org, and historical win/loss patterns by deal profile.

Authentication

The current demo deployment exposes endpoints without authentication. Production deployments should add an API key header or OAuth flow at the FastAPI middleware layer, scoped to read-only access on the deal endpoints.