




OneGrain

EVERY MARKET · BUILT READY

# Every Market, Ready.

The intelligence and execution layer for companies selling across fragmented markets, anywhere. It decides who to work today, micro-market by micro-market, then helps do the work and proves it in revenue.

OneGrain




# Markets are not linear. They fragment.

Wherever we look, the same shape repeats: one brand, one playbook, and underneath it hundreds of micro-markets that behave nothing alike. Different languages, seasons, clusters and calendars. The averages look fine. The opportunity hides in the variance.

Most companies respond by simplifying: fewer segments, broader campaigns, one national plan. The market does not simplify back.

**The truth:** fragmentation is not noise to average away. It is the map of where to act.



# This works wherever four conditions hold.

01



## Public data exists

Demographics, economy, seasonality and local signals are published and usable, so the outside world can be read, not guessed.

02



## Your data leaves trails

CRM, billing, app or transaction history that can be reconciled into one view of every account and customer.

03



## Channels can reach

WhatsApp, mail, phone or field teams that can act on a decision the same day it is made.

04



## The market fragments

Many micro-markets that one playbook cannot fit, so deciding locally beats broadcasting nationally.

## THE PROBLEM

# Revenue grows when headcount grows. That is the problem.

Most revenue organisations grow the same way: when the number has to rise, headcount rises. More field reps, more partners, more inside sales. Effort scales in a straight line, and so does its cost.

The result is arithmetic, not advantage. Each new rep carries roughly the same quota as the last. Each new territory costs what the previous one did. Acquisition cost climbs as teams push past the accounts they know, and pipeline keeps crowding into the same familiar places, because that is where the playbook was written.

Meanwhile the business already produces the raw material of better decisions: years of CRM history, billing records, inquiry logs, payment behaviour, wins and losses. It sits in systems that do not talk to each other, and it belongs to no one whose job is the only daily question that matters: who should we work today, and why?



*The constraint is not effort. It is that effort is allocated by habit, not by intelligence.*

# What your revenue stack already has. And what it is still missing.

## WHAT YOU HAVE

### In abundance.

- ✓ CRM with years of account history
- ✓ Field teams and channel partners
- ✓ Billing and transaction records
- ✓ Inquiry, messaging and call logs
- ✓ Quotation and win-loss data
- ✓ Campaign and event engagement

## WHAT IS MISSING



### The connective tissue.

- + **A micro-market-level view of demand** where the next order actually lives
- + **Prioritization tied to live signals** not last quarter's territory plan
- + **Timing intelligence** when an account is ready, not just whether
- + **Channel and language fit** visit, call or WhatsApp, in the buyer's words
- + **Closed-loop learning** every outcome makes the next decision sharper

# The intelligence layer above your revenue stack.

Not a CRM. Not a campaign tool. OneGrain sits above both: it decides where your effort should go, helps do the work, and learns from every outcome.

WHAT YOU GET

 **OneGrain**

**Decides** which account · why now · which channel · which language · what message  
**Executes** through your teams and channels, measures every outcome, and **learns** weekly

WHAT YOU ALREADY RUN



# We sell intelligence. Then we do the work.

Go-to-market is every decision about who to pursue, when, through which channel, and with what message. Most companies make those decisions by habit. OneGrain makes them with evidence, then acts on them.



## The intelligence

A daily, ranked answer for every territory: which accounts to work, why now, through which channel, in which language.



## The execution

We run the plays through the channels you already use: field visits, WhatsApp, calls, mail. Owned queues, with SLAs.



## The learning

Every outcome returns to the system. Decisions get sharper every week, while your headcount stays flat.

### TODAY'S QUEUE · ONE TERRITORY

#### Account #A-104

Reorder gap at 2.1x the local norm

WhatsApp · local language

#### Account #B-221

Quoted twice in nine days, still unsigned

Assisted call · today

#### Seasonal cohort

Buying window opens in this market this month

Timed offer · field

*Illustrative output. Identifiers stay hashed until dispatch.*

*A rep can work thirty accounts in a week. Intelligence decides which thirty.*

# Four layers. Built in sequence. Each layer earns the next.

01



## Reconcile

Normalize accounts across CRM, billing, distributor and inquiry records. Resolve duplicates and fragments. Attach every account to a micro-market, and every micro-market to its public context: demographics, industry clusters, seasonality, local calendars.

02



## See

One operating picture: pipeline by micro-market, cohort and product line. Dormancy, coverage gaps and whitespace made visible for the first time. Refreshed continuously, shared across sales, marketing and leadership.

03



## Decide

A daily ranked list per rep and per territory. Next best action, channel and timing for every account, with the reason stated. Rules carry the pilot; learned models take over as tracked outcomes accumulate.

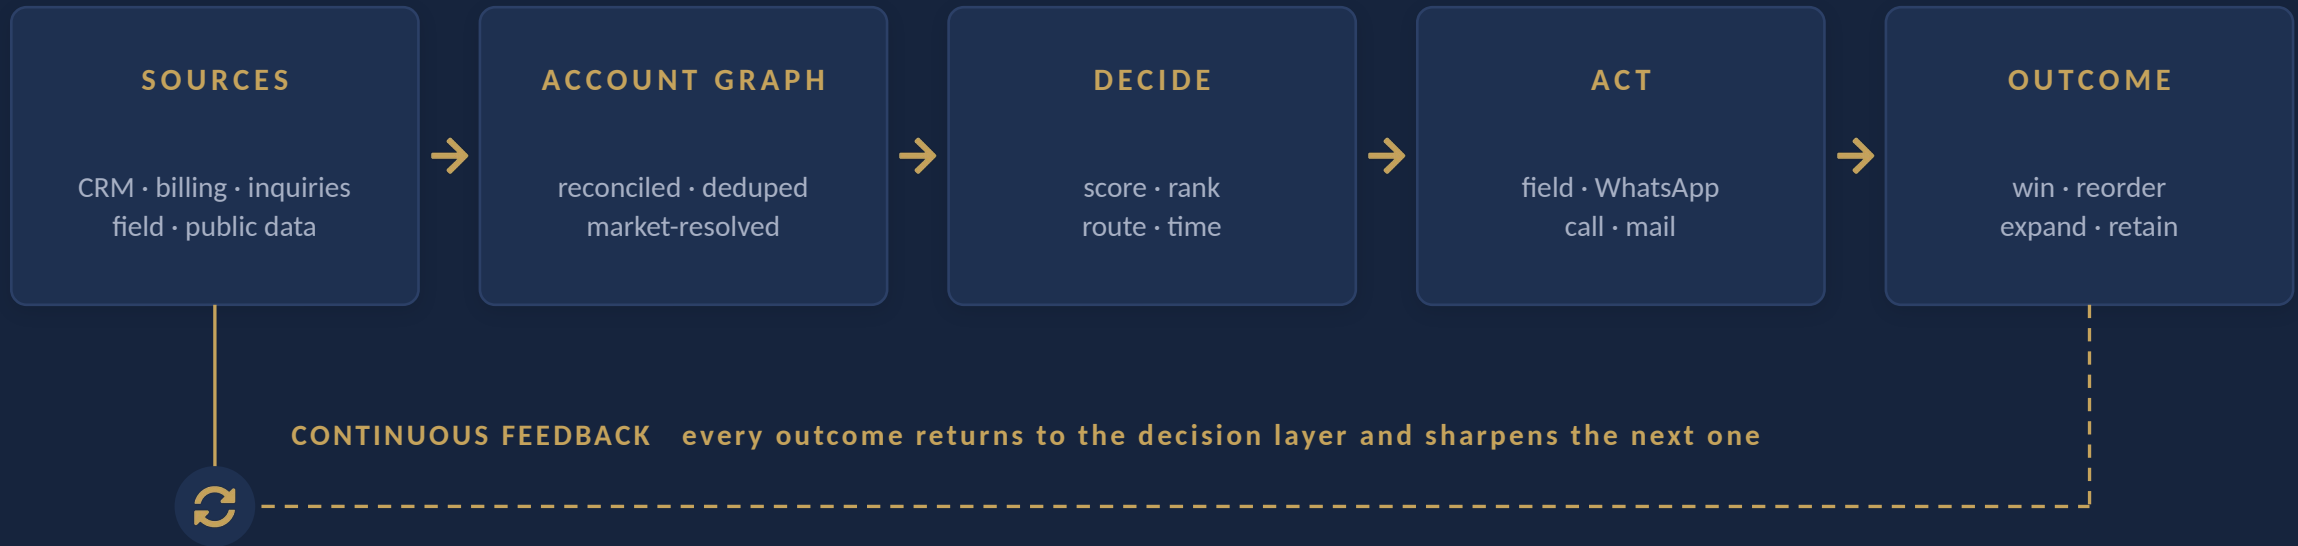
04



## Act

Queues with owners and SLAs. Interventions through field visits, WhatsApp, telecalling and mail. Closed-loop attribution from first touch to invoice, so every play is judged in revenue, not activity.

# How the system runs in production.



Public data sets the micro-market context. Your proprietary data makes it operational. **OneGrain's job is to make the two speak, and to point the result at revenue.**

# Two things compound. Neither is a feature.

Rules and routing can be rebuilt by anyone. These two cannot, because they get better with every market we run and every outcome we see.

## PILLAR ONE



### Public-data reconciliation

Every other tool reasons only over the data you already own, which is backward-looking and blind to the world outside. We fuse your internal records with a growing library of public context, so a score becomes a decision.



Demographics, density, clusters



Seasonality and local calendars



Weather, mobility, sentiment



Economic and vertical signals

## PILLAR TWO



### Fine-tuned GTM models

Beyond rules and general-purpose AI, we train models specifically for go-to-market: scoring, next-best-action, channel, timing and message, learned from real outcomes. Each cycle makes the next decision sharper.



Decision quality, cycle over cycle

# From a daily list to a layer of agents.

Today OneGrain hands your teams a ranked list and the reasoning behind it. The direction is a set of focused agents that carry each step of the loop, with your team approving the moves that matter. People stay in control; the busywork moves to the system.



## Signal agent

Watches behaviour and public data for the moment an account becomes worth a touch.



## Decision agent

Ranks who to work, picks the channel, language and timing, and states the reason.



## Outreach agent

Drafts the message per account and signal, inside templates your team approves.



## Routing agent

Sends self-serve where it fits, hands the rest to a rep with full context attached.



## Service agent

Catches failed payments, stalled onboarding and support gaps before they churn.



## Learning agent

Reads every outcome and retrains the models, so next week's decisions are sharper.

*Direction, not today's scope. Pilots run on the ranked list and your approvals. Agents come online as the loop earns the trust to automate each step.*

# Five trigger types. Tuned to the local calendar.



An account quiet past its local norm, queued for recovery before a competitor notices the silence.



Repeat inquiries, requotes, return visits. High intent flagged for assisted follow-up while it is still warm.



Fiscal cycles, festivals, seasons and local events, wherever the calendar moves demand, opened as timed plays per market.



Single-product accounts inside multi-product markets, surfaced for cross-sell with the offer that fits.



Shrinking orders, slowing payments, fading engagement. Caught while there is still time to act, not explained at quarter end.

**Each trigger generates a queue.**

Each queue has an owner, an SLA, and an outcome measured back to revenue. Nothing falls between teams.

# Beyond reporting: operating intelligence, by cohort and by play.



## COHORT INTELLIGENCE

**Which accounts move. Which never will.**

- By micro-market and cluster: who buys, who defers
- By size, credit behaviour and payment rhythm
- By language and channel responsiveness
- By season: when each micro-market opens its wallet
- By product mix and reorder cadence
- By source: which inquiries become invoices



## PLAY INTELLIGENCE

**Which moves work. On whom.**

- Which message converts in which micro-market
- Which channel earns a reply: visit, call or WhatsApp
- Which offers move which clusters
- Which sequences revive dormant accounts
- Which timing wins against which competitor
- Which plays to retire, and which to scale

# Illustrative ranges. Every one testable against your own data.

*Drawn from comparable enterprise engagements in distribution and adjacent verticals. These are starting hypotheses for the pilot, not promises.*

METRIC	TODAY (TYPICAL)	WITH ONEGRAIN	DELTA
Qualified meetings per rep, per week	Index 100	Index 130 – 160	+30 – 60%
Dormant-account reactivation	3 – 8%	12 – 20%	~2 – 3×
Cost per qualified opportunity	Index 100	Index 65 – 80	~20 – 35% lower
Time from signal to first touch	Days	Hours	10× faster
Cross-sell take-rate, existing accounts	2 – 4%	5 – 9%	~2× lift
Pipeline attribution	Partial	Closed loop	End to end

*Even a fraction of these moves, captured in one territory for one quarter, **pays for the engagement many times over**. That is the bet we are asking you to test.*

# Ninety days. One territory. One number at the end.

## WEEKS 1 - 3

### Reconcile

One territory instrumented end to end. Accounts resolved and placed in their micro-markets. Public context attached. Baseline established against your current run-rate, so the lift is measured, not narrated.

## WEEKS 4 - 7

### See & Decide

The operating picture stood up for your team. Dormancy, whitespace and coverage views live. Ranked daily queues running in shadow mode beside your current process, scored but not yet acting.

## WEEKS 8 - 12

### Act

Two trigger queues live in production: one reactivation, one timed micro-market play. Measured against a held-out control. Outcome reported in revenue terms, not in activity metrics.

## COMMERCIAL STRUCTURE

A bounded, paid pilot. Fixed fee, scoped to one territory and three trigger queues. The success number is agreed before we start, with a clear go / no-go decision at day 90.

# Three things. None of them heavy.



## Scoped data access

Read-only access to CRM, billing and inquiry data for the pilot territory. No production write access. Everything governed by your existing data-handling controls, inside your environment.



## One business owner

A single nominated owner on the revenue side who can decide which territory, which triggers and which plays to test. Empowered to act, not just kept informed.



## One technical contact

A nominated person from data or IT to sit alongside us for environment access, schema questions and deployment. Two hours a week, not a full-time secondment.

# Your data never meets our cloud. Non-negotiable, by design.



## Deployed in your infrastructure

Default deployment is on-premise or inside your own cloud. Models, pipelines and orchestration run where your data already lives. We bring the platform to the data, never the other way around.



## Hashed by design

Decisioning runs on hashed, pseudonymised identifiers. Names, numbers and identities resolve only inside your environment, at the moment of dispatch, by your systems.



## The strictest local law as the floor

GDPR, DPDP, PDPL or your local regime: we design to the strictest applicable standard, keep everything inside your environment, and exit with the playbook we earned and nothing else.

## THE ASK

# One meeting. One territory. Ninety days.

**01****Now**

A two-hour working session. We map where your funnel actually leaks, together, on a screen, with your own numbers in front of us.

**02****Week 2**

A scoped pilot proposal: one territory, three trigger queues, fixed fee, a clear measurement plan, and a go / no-go gate at day 90.

**03****Day 91**

A working demand engine, or an honest answer about why this was not the right time. Either way, you keep everything that was built.

*Wherever markets fragment, we are ready. One grain at a time.*